

Automating Verification of Functional Programs with Quantified Invariants

Abstract

We present the foundations of a verifier for higher-order functional programs with generics and recursive algebraic data types. Our verifier supports finding sound proofs and counterexamples even in the presence of certain quantified invariants and recursive functions. Our approach uses the same language to describe programs and invariants and uses semantic criteria for establishing termination. Our implementation makes effective use of SMT solvers by encoding first-class functions and quantifiers into a quantifier-free fragment of first-order logic with theories. We are able to specify properties of datastructure operations involving higher-order functions with minimal annotation overhead and verify them with a high degree of automation. Our system is also effective at reporting counterexamples, even in the presence of first-order quantification.

Categories and Subject Descriptors D.2.4 [Software Engineering]: Software/Program Verification; F.3.1 [Logics and Meaning of Programs]: Specifying and Verifying and Reasoning about Programs

Keywords software verification; higher-order functions; satisfiability modulo theories

1. Introduction

This paper presents the foundations of a system for verifying and falsifying (possibly quantified) properties of purely functional higher-order programs. The language we consider, although syntactically a subset of Scala, has simple type-theoretic foundations (Hindley-Milner type system with algebraic data types), but enriches it with specifications. Specifications include function pre- and postconditions as well as invariants, and they can refer to recursive functions and quantifiers over first-order values. Our tool supports automated finding of proofs and counterexamples for specifications and also establishes termination of functions to ensures soundness of reasoning. We illustrate our system on a set of examples.

Figure 1 shows the specification of a purely functional binary search tree storing unbounded integers (in source code denoted `BigInt` but here rendered \mathbb{Z} for brevity). The code shows the type definition as well as its functions `content`, `contains`, and `insert`. The data type invariant appears as a `require` clause in the constructor `Node`. While there are many ways to define invariants for simple binary search trees (including some that are executable), this one is closest to textbook description and it is very concise. The specification of methods in `ensuring` clauses also indicate that they behave as expected: `contains` corresponds to set membership and `insert` to insertion into a set. Our system verifies that all the specifications and invariants hold, that all pattern matching is complete, and that all functions terminate. The verification takes under 1.5 seconds, which is notably faster than executing the scalac compiler itself.

Figure 2 shows the use of functions stored inside case classes along with quantified invariants enforcing that stored functions satisfy the supplied conditions. Using such data type definition, it is possible to manipulate functions together with specifications and

```
sealed abstract class Tree {
  def content: Set[Z] = this match {
    case Leaf() => Set.empty[Z]
    case Node(l, v, r) => l.content ++ Set(v) ++ r.content
  }

  def insert(value: Z): Node = (this match {
    case Leaf() => Node(Leaf(), value, Leaf())
    case Node(l, v, r) =>
      if (v < value) Node(l, v, r insert value)
      else if (v > value) Node(l insert value, v, r)
      else Node(l, v, r)
  }) ensuring(_ .content == content ++ Set(value))

  def contains(value: Z): Boolean = (this match {
    case Leaf() => false
    case Node(l, v, r) =>
      if (v == value) true
      else if (v < value) r contains value
      else l contains value
  }) ensuring(_ == (content contains value))
}

case class Leaf() extends Tree
case class Node(left: Tree, value: Z, right: Tree) extends Tree {
  require(∀(x:Z) => (left.content contains x) => x < value) &&
    ∀(x:Z) => (right.content contains x) => value < x))
}
```

Figure 1. Binary Search Tree with invariant specified within the Node data type definition.

```
case class ~>[A,B](f: A => B, pre: A => Boolean, ens: B => Boolean) {
  require(∀(x: A) => pre(x) => ens(f(x)))

  def apply(x: A): B = {
    require(pre(x))
    f(x)
  } ensuring(ens) }

def map[A, B](list: List[A], f: A ~> B): List[B] = {
  require(∀(x:A) => list.contains(x) => f.pre(x))
  list match {
    case Cons(x, xs) => Cons[B](f(x), map(xs, f))
    case Nil() => Nil[B]()
  } } ensuring (res => ∀(x: B) => res.contains(x) => f.ens(x))
}
```

Figure 2. Encoding of first-class functions with contracts and its use to define a map on lists.

ensure that the conditions are propagated through data structures, as shown on the example of the map function on lists. The verification and termination of this program also succeed quickly.

Figure 3 shows a definition of (a relaxed form of) multisets using first-class functions. Using this data type, we specify merge sort, ensuring that the multiset of elements is preserved (Section 9 also discusses a version with generic elements mapped to keys.) Verification also succeeds quickly, and the only function for which the system does not automatically prove termination is merge itself,

```

case class Bag[T](f: T ⇒ Z) {
  def union(that: Bag[T]) = Bag((x: T) ⇒ f(x) + that.f(x))
  def equals(that: Bag[T]) = ∀(x: T) ⇒ f(x) == that.f(x) }

def content[T](list: List[T]): Bag[T] = list match {
  case Nil() ⇒ Bag(_ ⇒ 0)
  case Cons(x, xs) ⇒
    Bag(y ⇒ if (x == y) 1 else 0) union content(xs) }

def isSorted(list: List[Z]): Boolean = list match {
  case Cons(x1, tail @ Cons(x2, _)) ⇒ x1 ≤ x2 && isSorted(tail)
  case _ ⇒ true }

def merge(l1: List[Z], l2: List[Z]): List[Z] = {
  require(isSorted(l1) && isSorted(l2))
  (l1, l2) match {
    case (Cons(x, xs), Cons(y, ys)) ⇒
      if (x ≤ y) Cons(x, merge(xs, l2))
      else Cons(y, merge(l1, ys))
    case _ ⇒ l1 ++ l2 }
  } ensuring { res ⇒ isSorted(res) &&
    (content(res) equals (content(l1) union content(l2))) }

def split(list: List[Z]): (List[Z], List[Z]) = (list match {
  case Cons(x1, Cons(x2, xs)) ⇒
    val (s1, s2) = split(xs)
    (Cons(x1, s1), Cons(x2, s2))
  case _ ⇒ (list, Nil[Z]())
  }) ensuring { res ⇒
    (content(res._1) union content(res._2)) equals content(list) }

def mergeSort(list: List[Z]): List[Z] = (list match {
  case Cons(_, Cons(_, _)) ⇒
    val (s1, s2) = split(list)
    merge(mergeSort(s1), mergeSort(s2))
  case _ ⇒ list
  }) ensuring { res ⇒ isSorted(res) &&
    (content(res) equals content(list)) }

```

Figure 3. Higher-order functional encoding of multisets and the specification of mergeSort with multiset content semantics.

```

def associative[A](f: (A,A) ⇒ A): Boolean =
  ∀((x: A, y: A, z: A) ⇒ f(f(x, y), z) == f(x, f(y, z)))

def commutative[A](f: (A,A) ⇒ A): Boolean =
  ∀((x: A, y: A) ⇒ f(x, y) == f(y, x))

def rotates[A](f: (A,A) ⇒ A): Boolean =
  ∀((x: A, y: A, z: A) ⇒ f(f(x, y), z) == f(f(y, z), x))

def assocComm[A](f: (A,A) ⇒ A): Boolean = {
  require(associative(f))
  commutative(f) }.holds

def commAssoc[A](f: (A,A) ⇒ A): Boolean = {
  require(commutative(f))
  associative(f) }.holds

def commRotateAssoc[A](f: (A,A) ⇒ A): Boolean = {
  require(commutative(f) && rotates(f))
  associative(f) }.holds

```

Figure 4. Exploring dependencies between algebraic properties with counterexamples and proofs.

because the decrease in list size that split produces is not obvious to the system.

What makes our tool useful is that it can also report counterexamples. This feature works best for quantifier-free safety properties, but also works in some interesting cases for quantified properties. If in Figure 1 we swap the ordering in the appropriate line of contains to become **else if** ($v > \text{value}$) `r.contains(value)`, our system (given an appropriate command-line option) finds the counterexample; it also

confirms that the counterexample is real by performing execution that soundly interprets *a priori* unbounded quantifiers to confirm that the returned tree satisfies the invariants.

Our system can also find counterexamples for termination: for a function defining an interpreter for lambda calculus (Section 9) our termination checker ends up considering a possibly looping path and then uses the SMT solver to find that the well-known term $(\lambda x.xx)(\lambda x.xx)$ makes the path constraint satisfiable.

Figure 4 shows how function values and quantifiers enable abstraction of algebraic properties. Our system enables users to then reason about relationships between these properties. The shorthand holds is an abbreviation for **ensuring**(`_ == true`), and claims that a function with Boolean return type always returns **true**, so it can be used to encode conjectures and theorems. Our system finds counterexamples (over a domain of 7 uninterpreted elements) for the conjectures `assocCom` and `comAssoc` showing that neither of the properties associativity and commutativity implies the other. Furthermore, it proves `comRotateAssoc`, showing that commutativity and “rotation” imply associativity. These counterexamples and proofs are also found in time below 1.5 seconds.

Contributions. We present a verifier for higher-order functional programs with the ability to automatically find proofs and counterexamples. Our verifier incorporates the following new techniques:

- Verification of first-class functions along arbitrary paths in the program even when these are contained within input datastructures. We discuss a meaningful representation for reported counterexamples and present a sound and executable notion of first-class function equality. Our procedure is complete for counterexamples.
- Termination checking in the presence of higher-order functions, recursion and etc. We present a termination prover that builds upon the verification techniques we presented in a sound and efficient way, that conservatively allows the use of inductive reasoning in many useful cases.
- Support for arbitrarily nested ADT invariants that significantly reduce the annotation burden of proofs and provide better specification mechanisms for first-class functions. The technique we present is complete for counterexample finding as well.
- First-order quantifier instantiation with support for sound model finding. We present an extension to MBQI that enables model finding even in cases where the clause set is not *stratified* and discuss certain decidability implications.
- We discuss how the different presented techniques integrate with each other and consider soundness of the unified procedure.

We have evaluated our verifier on a set of benchmarks covering a wide range of features.

2. Preliminaries

Defining the programs. We start by defining the purely functional subset of Scala (called PureScala) on which our transformation is defined. A PureScala program P consists of a set of *definition* non-terminal expansions as well as an *expr* expansion. An expression can be a primitive operation, conditional, pattern matching, application of named or of first-class function. Algebraic data types are written as simple forms of Scala case classes. (Figure 6 in the Appendix shows the syntax.) The set of definitions that are relevant to the given expression is defined by the expression itself and is therefore not explicitly considered hereafter. Note that we use the terms program and expression interchangeably for this reason.

The typing rules for PureScala are as expected for call by value simply typed lambda calculus with recursive sum and product types (called algebraic data types or ADTs hereafter). We generally

assume that considered programs and expressions are well-typed. We therefore associate to each expression E its type $E.tpe$. We implicitly assume that variables and first-class function definitions are each associated with some fixed type, providing the initial Γ context for $E : E.tpe$.

Given an algebraic data type (ADT) T , we associate to T a set of constructors $\{\text{cons}_{T,1}, \dots, \text{cons}_{T,n}\}$ such that each $\text{cons}_{T,i}$ is a m_i -ary function of type $T_{i,1} \times \dots \times T_{i,m_i} \rightarrow T$. Instances of type T therefore have the shape $\text{cons}_{T,i}(E_1, \dots, E_{m_i})$. We omit the T index and write cons_i when the relevant ADT type is clear from context. Expressions within an ADT instance (*i.e.* ADT *fields*) are accessed through pattern match case deconstruction.

We do not distinguish between recursive and anonymous function definitions (unless explicitly stated otherwise) as all functions are considered first-class. We use f, g, \dots to refer to first-class function expressions. We further associate to function f the following variable sequences $f.args$ corresponding to its arguments and $f.closures$ corresponding to some deterministic ordering of the free variables in f . Finally, we have the expression $f.body$ that corresponds to the body of f .

We define $\text{vars}(P)$ to be the *free variables* within program/expression P . If an expression contains no free variables, we say it is *ground*. We define the notion of *value* as

$$\text{value} ::= \text{true} \mid \text{false} \mid \text{cons}_i(\text{value}, \dots, \text{value}) \mid \text{ground function } f$$

Given a program P , we call a mapping m that relates each variable $v \in \text{vars}(P)$ to some value $V : v.tpe$ a set of *inputs* to P . Note that $m \llbracket P \rrbracket$ must be ground, and evaluation semantics follow those of the simply typed lambda calculus with conditionals. Further note that any well-typed program P with inputs m will either evaluate to some value or diverge, evaluation cannot get *stuck*.

Basic formula terminology. A *signature* Σ consists of a set of *sorts*, one of *function* symbols, and a set of *predicate* symbols. We view predicates as functions with boolean-sorted result. Each function symbol is associated with a sequence of argument sorts (whose length is the function's *arity*) and a result sort. Functions with zero arity are called *constants*. We let $\text{vars}(\phi)$ consist of the set of uninterpreted constants in ϕ . A *term* can be either a constant or a function application $f(t_1, \dots, t_n)$ where f is a function symbol and t_1 to t_n are terms. Boolean-sorted terms are sometimes called *atoms*, and we generically use *clause* to refer to any production using atoms, disjunction (\vee), conjunction (\wedge), and their derivatives as these can trivially be translated into CNF form. We sometimes use set inclusion notation over finite sets to denote equality disjuncts. We are generally interested in satisfiability of clause sets where uninterpreted function symbols are taken as existentially quantified.

We are also interested in Σ -theories, namely sets of Σ -sentences that are closed under logical deduction. We assume each theory T is associated to a single sort σ_T , and can introduce a set of interpreted function symbols whose interpretations are restricted by the models for T . We use the \simeq symbol for the interpreted equality symbol. We are specifically interested in the theory of algebraic data types where a specific ADT theory T is defined similarly to an ADT T in PureScala. Namely, we associate to T the following interpreted function symbols for $1 \leq i \leq n$: $\text{isCons}_{T,i}$, $\text{cons}_{T,i}$, and $\text{field}_{T,i,j}$ for $1 \leq j \leq m_i$, as well as the usual axiomatization. We again omit the T index when it is clear from context. We assume theories are *separate*. Namely, given a theory T with associated sort σ_T , given an interpreted function symbol f associated to T and the following interpretation $t = M \llbracket f(t_1, \dots, t_n) \rrbracket$, either i) the sort of t is σ_T , or ii) $t \subseteq M \llbracket t_i \rrbracket$ for some $i : 1 \leq i \leq n$. Note that this assumption does not hold for all interesting theories supported by SMT solvers (*e.g.* consider sets with cardinality), but it does hold for the theories on which our transformation depends. Moreover, our

definition of *theory separation* can be further augmented to allow t with *non-container* sort such as integer or real. Such considerations are out of the scope of this paper.

Finally, we say a term t is *ground* iff it contains no uninterpreted symbol. Given a ground term t and the equivalence class of ground terms $G_t = \{t_2 \mid t \simeq t_2\}$, we call the term $t_{min} \in G_t$ with smallest size a *value* (given multiple terms with same size, one is selected by convention). Term interpretation in a model is assumed to always produce a value.

3. Verification of Programs with Function Values

In this section we introduce mapping from PureScala program expressions (containing first-class function applications) to quantifier-free terms and formulas suitable for an SMT solver that has no knowledge of recursive or first-class functions. We build on this translation to present a procedure to prove that a boolean-result PureScala expressions cannot return false. Our procedure iteratively refines information about named and anonymous functions, generating a sequence of quantifier-free formulas that constitute increasingly precise program approximations.

3.1 Transformation from Expressions to Formulas

The main difference between source expressions and target formula terms is that functions in the target formulas are represented using values from an uninterpreted domain. The translation function explicitly represents increasing amount of information about the applications of such functions.

From types to sorts. We first introduce a translation function \mathcal{S} from PureScala types T to sorts. When T is boolean or an algebraic data type, then $\mathcal{S}(T)$ is an appropriately chosen unique formula sort. For function types, however, we introduce a special sort σ_f with infinite cardinality. All function types map to σ_f . The only relation symbol defined on σ_f is \simeq and represents a notion of equality.

Our typing relation supports only invariant type parameters. We can therefore consider that instantiations of an ADT type T with different type arguments correspond to disjoint sorts. This consideration further extends to generic function definitions. We can thus ignore generic type parameters in the following by assuming (potentially infinite) sets of fully typed definitions. As these sets are explored incrementally through as-needed type substitutions in the access of $f.body$, the relevant portions of the definition sets remain finite. Finally, we allow uninstantiated type parameters in our formulas and let \mathcal{S} map to some (unique) sort with infinite cardinality and equality (cardinality limitations are outside the scope of this paper).

Program names to formula constants. We define an injection \mathcal{V} from PureScala variables to formula constants. (Such constants are, effectively, existentially quantified because we will be checking satisfiability of generated formulas.) For a variable $v : T$ the result of mapping $\mathcal{V}(v)$ has sort $\mathcal{S}(T)$.

Function applications. Our procedure relies on tracking two sets that are computed by the transformation. The first set, A , tracks function applications and consists of 4-tuples $(b, c, tpe, [a_1, \dots, a_n])$ where b is a Σ -constant specifying the *path condition* under which this function application can take place, c is a Σ -term with sort σ_f specifying the caller, tpe is the PureScala function type corresponding to c , and $[a_1, \dots, a_n]$ are Σ -terms argument to c (sometimes written a^n for brevity). We define the projection functions $(\cdot)_b$, $(\cdot)_c$ and $(\cdot)_{tpe}$ for $app \in A$. The second set, F , tracks $(b, f, \mathbf{f}, [c_1, \dots, c_n])$ 4-tuples where b is a Σ -constant as in A , f is a (fresh) Σ -constant with sort σ_f , \mathbf{f} is a first-class PureScala function discovered during translation, and c^n are Σ -terms corresponding to

the closures of \mathbf{f} . We also define the projection functions $(\cdot)_b$, $(\cdot)_f$ and $(\cdot)_{\mathbf{f}}$ for $fun \in F$.

Transformation rules. The transformation $(b, E) \rightsquigarrow (e, \phi, A, F)$ takes as input the Σ -constant b describing the current path condition and the expression E in PureScala and returns a 4-tuple where e is a Σ -term, ϕ is a set of clauses, and A, F are defined above. We further define the following projective transformation $(b, E) \rightsquigarrow^t e$. The transformation is defined inductively, so given $\{\text{SUB}_1, \dots, \text{SUB}_n\}$ the set of all direct children of E (expression children can be inferred from the definition of PureScala's syntax in Figure 6) with $(b, \text{SUB}_i) \rightsquigarrow (_ , \phi_i, A_i, F_i)$, we generally have $\phi = \bigcup_i \phi_i$, $A = \bigcup_i A_i$ and $F = \bigcup_i F_i$. Therefore, it is useful for the sake of clarity to introduce the difference projection $(b, E) \rightsquigarrow^t \phi - \bigcup_i \phi_i$. Note that the result of $\rightsquigarrow^t \phi$ completely defines the ϕ resulting from the transformation \rightsquigarrow of (b, E) (and vice-versa). We similarly introduce \rightsquigarrow^A and \rightsquigarrow^F . The rules of inference described below assume by convention that $(b, E) \rightsquigarrow^t \emptyset$, $(b, E) \rightsquigarrow^A \emptyset$, and $(b, E) \rightsquigarrow^F \emptyset$ when left unspecified.

We start by defining the fairly straightforward transformation of PureScala expressions for which closely corresponding encodings exist in the formula domain.

$$(b, v) \rightsquigarrow^t \mathcal{V}(v) \quad (b, \text{true}) \rightsquigarrow^t \text{true} \quad (b, \text{false}) \rightsquigarrow^t \text{false}$$

$$\frac{(b, E_1) \rightsquigarrow^t e_1 \quad \dots \quad (b, E_m) \rightsquigarrow^t e_m}{(b, \text{cons}_i(E_1, \dots, E_m)) \rightsquigarrow^t \text{Cons}_i(e_1, \dots, e_m)}$$

In order to handle general function definitions without resorting to quantifiers, we translate these to constants and perform incremental dynamic dispatch over call sites discovered during the transformation. This leads to the following rules for function handling

$$\text{FUN} \frac{f \text{ fresh with sort } \sigma_{\mathbf{f}} \quad \mathbf{f} \text{ function}}{(b, \mathbf{f}) \rightsquigarrow^t f \quad (b, \mathbf{f}) \rightsquigarrow^A \mathbf{f} \{ (b, f, \mathbf{f}, \mathcal{V}(\mathbf{f}. \text{closures})) \}}$$

$$\text{APP} \frac{(b, c) \rightsquigarrow^t c \quad (b, E_1) \rightsquigarrow^t e_1 \quad \dots \quad (b, E_n) \rightsquigarrow^t e_n}{(b, C(E_1, \dots, E_n)) \rightsquigarrow^t \text{dispatch}_{c.tpe}(c, e_1, \dots, e_n) \quad (b, C(E_1, \dots, E_n)) \rightsquigarrow^A \mathbf{f} \{ (b, c, C.tpe, [e_1, \dots, e_n]) \}}$$

where the FUN rule accumulates potential dispatchees, and the APP rule both introduces well-typed dispatch calls through type-indexed function symbols and perform the necessary bookkeeping for future inlining.

We now define the transformation of **if**-expressions and discuss how they introduce sound path conditions. The transformation of a (b, E) pair where

$$E ::= \mathbf{if} (\text{COND}) \text{ THEN } \mathbf{else} \text{ ELSE}$$

relies on two boolean-sorted fresh constants b_t and b_e , one for each new path condition (namely COND and \neg COND). We also need a fresh constant r with sort $\mathcal{S}(\top)$ where $E : \top$. Given $(b, \text{COND}) \rightsquigarrow^t c$, $(b_t, \text{THEN}) \rightsquigarrow^t e$, and $(b_e, \text{ELSE}) \rightsquigarrow^t e$, we can define the clausal encoding of the **if**-expression as

$$\phi_{branch} = \{ (b \wedge c) \iff b_t, \quad b_t \implies (r \simeq t), \\ (b \wedge \neg c) \iff b_e, \quad b_e \implies (r \simeq e) \}.$$

Note that computation of t and e relies on the constants b_t and b_e , such that the clauses (and other bookkeeping information) resulting from the transformation of THEN and ELSE depend respectively on b_t and b_e . Furthermore, the clauses in ϕ_{branch} ensure that the data-flow can only take a single branch of the conditional, thus ensuring b_t and b_e are sound path conditions. The definition of

ϕ_{branch} naturally leads to the following inference rule

$$\text{IF} \frac{b_t, b_e, r \text{ fresh} \quad (b, \text{COND}) \rightsquigarrow (c, \phi_c, A_c, F_c) \quad (b, \text{THEN}) \rightsquigarrow (t, \phi_t, A_t, F_t) \quad (b, \text{ELSE}) \rightsquigarrow (e, \phi_e, A_e, F_e)}{(b, E) \rightsquigarrow (r, \phi_c \cup \phi_t \cup \phi_e \cup \phi_{branch}, A_c \cup A_t \cup A_e, F_c \cup F_t \cup F_e)}$$

Finally, encoding of pattern match expressions is similar to that of **if**-expressions and can be found in the appendix.

Properties of transformation. We will now state several relevant properties of the transformation \rightsquigarrow . In this next part, we will consider the expression E with inputs m such that $m \llbracket E \rrbracket \xrightarrow{*} v$. We further consider $(b, E) \rightsquigarrow (e, \phi_e, A_e, F_e)$ and $(b, v) \rightsquigarrow (v, \phi_v, _ , _)$ as given. We start by noting that the transformation \rightsquigarrow corresponds to an under-approximation of the expression E with respect to evaluation semantics.

Lemma 1. *There exists model $M \models b \wedge \phi_e \wedge \phi_v \wedge e \simeq v$.*

Consider $E_i \subseteq E$ such that $(b_i, E_i) \rightsquigarrow^t t_i$ occurred during the transformation of E . Note that all $E_i \subseteq E$ that are not contained within an anonymous function definition are visited by \rightsquigarrow . The position of E_i uniquely defines t_i , and vice-versa. Furthermore, they each uniquely define b_i , leading to the well-founded definitions $C_t(E_i) = t_i$, $C_E(t_i) = E_i$, $C_b(E_i) = b_i$, and $C_b(t_i) = b_i$. Furthermore, if the position E_i is encountered during evaluation of $m \llbracket E \rrbracket$, there exists a unique value E_v corresponding to that position. We can therefore define $C_v(E_i) = E_v$ if E_i is encountered during evaluation, and $C_v(E_i) = \perp$ otherwise.

Proposition 2. *For $M \models b \wedge \phi_e \wedge \phi_v \wedge e \simeq v$ and $E_i \subseteq E$, $M \models C_b(E_i)$ iff $C_v(E_i) \neq \perp$.*

As our encoding of first-class functions introduces fresh constants, the set of models satisfying Lemma 1 will be underconstrained in general. We therefore introduce the following lemma that deals more specifically with these constants

Lemma 3. *There exists an M satisfying Lemma 1 such that for $E_c \subseteq E$ where $E_c ::= C(E_1, \dots, E_n)$ and for $fun \in F_e$, if $M \models C_b(E_c) \wedge (fun)_b$, then $M \models C_t(C) \simeq (fun)_f$ iff $C_v(C) = C_v(C_e((fun)_f))$.*

3.2 Unfolding Function Applications

We discuss how our procedure relates the uninterpreted function applications introduced by the APP rule to the associated function definitions.

We start by considering a constant-program pair (b_p, P) with $(b_p, P) \rightsquigarrow (e_p, \phi_p, A_p, F_p)$. Given $app \in A_p$ and $fun \in F_p$ where $app = (b_c, c, tpe, a^n)$, $fun = (b_f, f, \mathbf{f}, c^n)$, and $tpe = \mathbf{f}.tpe$, consider the transformation $(b_{\mathbf{f}}, \mathbf{f}.body) \rightsquigarrow (e_{\mathbf{f}}, \phi_{\mathbf{f}}, A_{\mathbf{f}}, F_{\mathbf{f}})$ where $b_{\mathbf{f}}$ is a fresh constant. Further consider the substitution

$$\theta_{\mathbf{f}} = \{ \mathcal{V}(\mathbf{f}.args) \rightarrow a^n \} \{ \mathcal{V}(\mathbf{f}.closures) \rightarrow c^n \}.$$

Note that the first portion of $\theta_{\mathbf{f}}$ is entailed by the evaluation rules on function applications and the second part corresponds to the semantics of $m \llbracket P \rrbracket$ for inputs m . Now, consider the clause

$$\text{dispatch}_{\mathbf{f}.tpe}(c, a_1, \dots, a_n) \simeq \theta_{\mathbf{f}}[e_{\mathbf{f}}]. \quad (\text{disp})$$

Under the conditions b_c, b_f and $c \simeq f$, the (disp) clause along with the substituted sets $\theta_{\mathbf{f}}[\phi_{\mathbf{f}}]$, $\theta_{\mathbf{f}}[A_{\mathbf{f}}]$, and $\theta_{\mathbf{f}}[F_{\mathbf{f}}]$ corresponds to taking the transformation *after* having inlined the function call associated with app in P . We therefore define

$$\mathcal{I}(app, fun) = (\theta_{\mathbf{f}}[\phi_{\mathbf{f}}] \cup \{ b_{\mathbf{f}} \iff (b_c \wedge b_f \wedge c \simeq f), \\ b_{\mathbf{f}} \implies (\text{disp}) \}, \theta_{\mathbf{f}}[A_{\mathbf{f}}], \theta_{\mathbf{f}}[F_{\mathbf{f}}])$$

which enables us to progressively refine the transformation of P by introducing new clauses corresponding to inlinings of function applications in P . Note that we preserve the $b \implies c$ structure of our clause set. Furthermore, if fun doesn't correspond to C , then all new elements will be irrelevant as $b_{\mathbf{f}}$ will not hold.

Lemma 4. *Given inputs m with $m[\![P]\!] \xrightarrow{*} \vee$ and $(b_p, \vee) \rightsquigarrow^t v_p$, given $E_c \subseteq P$ where $C_t(E_c) = \text{dispatch}_{tpe}(c, a_1, \dots, a_n)$ and $(C_b(E_c), C_v(E_c)) \rightsquigarrow^t v_c$, there exists model M such that either*

1. $M \models b_p \wedge \phi_p \wedge e_p \simeq v_p \wedge \theta_{\mathbf{f}}[\![\phi_{\mathbf{f}}]\!] \wedge \theta_{\mathbf{f}}[\![e_{\mathbf{f}}]\!] \simeq v_c$, or
2. $M \models b_p \wedge \phi_p \wedge e_p \simeq v_p \wedge \neg b_{\mathbf{f}}$.

Based on these considerations and assuming $P : \text{Boolean}$, we define the sequence $U(P) = u_0, u_1, \dots$ of triplets $u_i = (\phi_i, A_i, F_i)$ where $u_0 = (\phi_p \cup \{b_p, \neg e_p\}, A_p, F_p)$ and u_i is defined given $app_i \in A_{i-1}$ and $fun_i \in F_{i-1}$ as $u_i = u_{i-1} \cup \mathcal{I}(app_i, fun_i)$. The sequence of u_i can therefore be seen as progressively inlining function applications inside the program. This definition leads to the following theorem

Theorem 5. *Given $P : \text{Boolean}$ with ϕ_i from $U(P)$, if $\phi_i \in \text{Unsat}$, then no inputs m exist such that $m[\![P]\!] \xrightarrow{*} \text{false}$.*

Given some expression $E : \text{Boolean}$, if there exist no inputs m such that $m[\![E]\!] \xrightarrow{*} \text{false}$, we say E is *valid*. We are generally interested in this work in either showing validity (as is the case for Theorem 5), or reporting inputs for which the given expression evaluates to **false**.

4. Inductive Assumptions Modulo Termination

In practice, most interesting properties dealing with recursive functions will never have $\phi_i \in \text{Unsat}$ as the uninterpreted dispatch symbols offer too high a degree of freedom. It therefore becomes useful to introduce some notion of inductiveness into the procedure. Consider an expression $\text{POST} : \text{Boolean}$ such that there exists a function application $\mathbf{f}(v_1, \dots, v_n) \subseteq \text{POST}$ where $\mathbf{f}.closures = \emptyset$, $\text{vars}(\text{POST}) = \{v_1, \dots, v_n\}$, and v_1, \dots, v_n are taken as quantified. We extend the procedure to produce sound proofs in the presence of assumption POST when \mathbf{f} is terminating.

Given app_i, fun_i selected in $U(P)$ where $(fun_i)_{\mathbf{f}} = \mathbf{f}$, consider $b_{\mathbf{f}}$ and $\theta_{\mathbf{f}}$ introduced by $\mathcal{I}(app_i, fun_i)$. Recall that $b_{\mathbf{f}}$ holds iff $(app_i)_b \wedge (fun_i)_b \wedge (app_i)_c \simeq (fun_i)_f$ and $\theta_{\mathbf{f}} = \{\mathcal{V}(\mathbf{f}.args) \mapsto a^n\}$ for a^n from app_i (the second part of $\theta_{\mathbf{f}}$ is omitted as $\mathbf{f}.closures = \emptyset$). Now consider the transformation $(b_{\mathbf{f}}, \text{POST}) \rightsquigarrow (p, \phi_p, A_p, F_p)$ and let

$$\mathcal{P}_{\text{POST}}(app_i, fun_i) = (\theta_{\mathbf{f}}[\![\phi_p]\!] \cup \{b_{\mathbf{f}} \implies \theta_{\mathbf{f}}[\![p]\!]\}, \theta_{\mathbf{f}}[\![A_p]\!], \theta_{\mathbf{f}}[\![F_p]\!]).$$

For simplicity, we extend $\mathcal{P}_{\text{POST}}$ to a total function by letting $\mathcal{P}_{\text{POST}}(app_i, fun_i) = (\emptyset, \emptyset, \emptyset)$ when $(fun_i)_{\mathbf{f}} \neq \mathbf{f}$. We then also extend the procedure $U(P)$ to $U_{\text{POST}}(P)$ such that u_i is now defined as $u_i = u_{i-1} \cup \mathcal{I}(app_i, fun_i) \cup \mathcal{P}_{\text{POST}}(app_i, fun_i)$.

We want to show that the result of Theorem 5 is preserved in the extended $U_{\text{POST}}(P)$. Clearly, if there exist inputs m to POST such that $m[\![\text{POST}]\!] \xrightarrow{*} \text{false}$, then Lemma 1 tells us that there exists $M \models b_{\mathbf{f}} \wedge \phi_p \wedge \neg p$. Consequently, there exists an app, fun pair where $\mathcal{P}_{\text{POST}}(app, fun) = (\phi, _, _)$ such that $\phi \in \text{Unsat}$. As the existence of inputs m to POST does not imply the inexistence of inputs m_p to P such that $m_p[\![P]\!] \xrightarrow{*} \text{false}$, such a case leads to unsoundness. In order to guard against this occurrence, we must therefore show that there exist no such inputs m .

If we consider $U_{\text{POST}}(\text{POST})$ directly, our clause set ϕ_i will, at some point, contain both $\neg p$ from the initial transformation of POST and p from $\mathcal{P}_{\text{POST}}(app_i, fun_i)$, thus leading to $\phi \in \text{Unsat}$ even when inputs exist. We therefore consider the expression E_{POST}

defined as

$$E_{\text{BODY}} ::= \{\mathbf{f}.args \mapsto v^n\}[\![\mathbf{f}.body]\!] \\ E_{\text{POST}} ::= \{\mathbf{f}(v_1, \dots, v_n) \mapsto E_{\text{BODY}}\}[\![\text{POST}]\!]$$

Lemma 6. *Given ϕ_i from $U_{\text{POST}}(E_{\text{POST}})$, if $\phi_i \in \text{Unsat}$ and POST terminates on all inputs, then there exist no inputs m such that $m[\![\text{POST}]\!] \xrightarrow{*} \text{false}$.*

Showing validity of POST is interesting in its own right, however it is can be useful to consider arbitrary property P .

Theorem 7. *Given ϕ_i from $U_{\text{POST}}(P)$, if $\phi_i \in \text{Unsat}$, POST is valid, and POST terminate on all inputs, then there exist no inputs m such that $m[\![P]\!] \xrightarrow{*} \text{false}$.*

Note that we can trivially extend the $U_{\text{POST}}(\cdot)$ procedure to a set of expressions $\text{POST}_1, \dots, \text{POST}_m$ as all $\mathcal{P}_{\text{POST}_i}$ are independent. However, the termination condition on each POST_i becomes somewhat more complex in this setting. Indeed, divergence may appear due to inter-dependencies between the different assumptions even though each POST_i terminates for all inputs on its own. We associate to each POST_i its relevant function application $\mathbf{f}_i(v_{i,1}, \dots, v_{i,n_i})$. Given inputs m , we inductively define the set S_m of all expressions relevant to termination

1. $m[\![\text{POST}_i]\!] \in S_m$,
2. if $\mathbf{f}_i(E_1, \dots, E_n)$ was seen during evaluation of $s \in S_m$, then $\{v_k \mapsto v_k \mid 1 \leq k \leq n\}[\![\text{POST}_i]\!] \in S_m$.

We say $\text{POST}_1, \dots, \text{POST}_m$ are *post-terminating* for inputs m iff for all $s \in S_m$, evaluation of s is terminating. Note that for a single POST , post-termination for all inputs is equivalent to termination for all inputs. Let us now define $U_{\text{POST}}(\{\text{POST}_1, \dots, \text{POST}_m\}, P)$ in the obvious way using the relevant $\mathcal{P}_{\text{POST}_i}$ in the inductive definition of u_i . We can thus conclude with the following corollary

Corollary 8. *Given ϕ_i from $U_{\text{POST}}(\{\text{POST}_1, \dots, \text{POST}_m\}, P)$, if $\phi_i \in \text{Unsat}$, $\text{POST}_1, \dots, \text{POST}_m$ are valid, and $\text{POST}_1, \dots, \text{POST}_m$ are post-terminating for all inputs, then there exist no inputs m such that $m[\![P]\!] \xrightarrow{*} \text{true}$.*

5. Proving Termination

The soundness for inductive assumptions discussed in the previous section relies on termination of the expression for all inputs. Our verification framework therefore features an automated termination checker, which checks termination of each function definition. In principle, we could choose any method to ensure termination of computations, as long as we only use inductive reasoning for computations we have shown terminating. Note that it is also permissible to use properties previously proven by induction to establish termination of subsequent properties, which allows proving termination using non-trivial semantic arguments. Our process of verifying a program thus involves an interleaved sequence of verification and termination checks, as is common in interactive proof assistants.

Our termination checker requires that the recursion in type definitions is restricted to recursive reference of a type (with the same formal type parameters) in the field position. In particular, it disallows recursive definitions in which the type-level recursion involves function type constructor. The data types that the termination checker accepts can therefore be encoded in polymorphic lambda calculus [31], [28, Chapter 23]. Furthermore, function values stored in the arguments of the function whose termination is being checked are assumed to be terminating.

Given our definition of evaluation semantics, an infinite sequence of evaluation steps must contain an infinite subsequence of function application reduction steps. In the following, we distinguish between

anonymous functions of the shape $(x_1, \dots, x_n) \Rightarrow E$ (which are considered first-class) and named functions (which are not). We assume named functions are lifted to anonymous functions when appearing in first-class positions.

It is useful when constructing termination proofs to be able to consider some notion of *path* under which a given sub-expression can be evaluated. Namely, given expressions $E, S \subseteq E$ and PATH , if for all inputs m we have $m[\text{PATH}] \xrightarrow{*} \text{true}$ iff $m[E]$ encounters the position of S during evaluation, then we call PATH the *path* to S in E . Given our formalization of PureScala, it is difficult to consider the notion of *path* in the presence of pattern match expressions. Let us therefore consider the extension to PureScala with the following expression non-terminals $\text{isCons}_{\tau,i}(\cdot)$ and $\text{field}_{\tau,i,j}(\cdot)$, as well as the corresponding trivial extensions to \rightsquigarrow . It is clear that given a program P , we can now transform it into an equivalent program (with respect to evaluation and \rightsquigarrow transformation) that contains no pattern match expressions. We assume hereafter that all considered expressions have been thus transformed. We can now talk about the path under which an expression $E \subseteq P$ will be evaluated (noted $\text{path}(E)$), computed as the conjunction of all if-expression conditionals (respectively their negation depending on which branch contains E) that must hold for evaluation to reach E . For $E \subseteq ((x_1, \dots, x_n) \Rightarrow R) \subseteq P$, we relax our definition of *path* to avoid tracking the anonymous function to its application points and let $\text{path}(E)$ contain the free variables x_1, \dots, x_n . Note that when these are taken as universally quantified, this constitutes a sound approximation of first-class function application points.

Our aim is to show that all evaluation traces must be finite by establishing that the set of named function application reduction steps in each trace must be finite. Given the expression E , we therefore consider the set

$$C(E) = \{ \mathbf{f}(A_1, \dots, A_n) \mid \mathbf{f}(A_1, \dots, A_n) \subseteq E \}$$

and for each named function \mathbf{f} in P , we let $C_{\mathbf{f}} = C(\mathbf{f}.body)$. As in the path case, given anonymous function $(x_1, \dots, x_n) \Rightarrow E$ defined within $\mathbf{f}.body$, all named function applications within E belong to $C_{\mathbf{f}}$ as well. We then define the following relation

$$\frac{\begin{array}{l} A_1, \dots, A_n \text{ values} \quad m_{free} \text{ inputs} \quad \mathbf{g}(E_1, \dots, E_m) \in C_{\mathbf{f}} \\ \theta = m_{free} \cup \{ v_k \mapsto A_k \mid v_k \in \mathbf{f}.args \} \\ \theta[\text{path}(\mathbf{g}(E_1, \dots, E_m))] \xrightarrow{*} \text{true} \quad \theta[E_i] \xrightarrow{*} v_i, 1 \leq i \leq m \\ \mathbf{f}(A_1, \dots, A_n), \mathbf{g}(v_1, \dots, v_m) \in R \end{array}}{\theta \xrightarrow{*} \text{true}}$$

Note that we use the extra inputs m_{free} to cover the free variables in $\mathbf{g}(E_1, \dots, E_m)$ that appear due to anonymous function arguments.

We assume without loss of generality that $P ::= \mathbf{f}_0(v_1, \dots, v_n)$ as any program P can be brought into such a form by introducing a fresh named function \mathbf{f}_0 such that $\mathbf{f}_0.body = P$ and $\mathbf{f}_0.args = \text{vars}(P)$. Note that as \mathbf{f}_0 is fresh, it will only appear once in any evaluation trace given inputs. Now for inputs m , consider an evaluation trace such that the named function application $\mathbf{g}(B_1, \dots, B_m)$ is due to be reduced at some point in the trace, $\mathbf{g} \neq \mathbf{f}_0$ and B_1, \dots, B_m are values. As evaluation has reached this point, there must exist some named function application reduction $\mathbf{f}(A_1, \dots, A_n)$ in the trace such that $(\mathbf{f}(A_1, \dots, A_n), \mathbf{g}(B_1, \dots, B_m)) \in R$. Indeed, even when $\mathbf{g}(B_1, \dots, B_m)$ is contained within a closure which has arbitrarily travelled through the program, this closure must have been created at some point in the body of some $\mathbf{f}(A_1, \dots, A_n)$. Furthermore, as we universally quantify over potential closure arguments, any concrete arguments given to the closure will have been considered during construction of R , hence the inclusion. We assume that input first-class functions are terminating and support the assumption by checking that all functions defined in program terminate. This leads to the following statement

Proposition 9. *Given inputs m , if the trace of $m[\mathbf{f}_0.body]$ contains infinitely many named function application reduction steps, then there must exist an infinite sequence $\text{APP}_0, \text{APP}_1, \dots$ such that $\text{APP}_0 ::= \mathbf{f}_0(m[\mathbf{f}_0.args])$ and $(\text{APP}_i, \text{APP}_{i+1}) \in R$.*

A well-understood and general technique used to show finiteness of named function application chains consists of selecting a *measure* from the domain of a well-founded relation and showing that it decreases in all possible chains. Let us consider the following measure that corresponds to the structural size of the provided argument. For each ADT type T in the program, let

$$\text{def size}_{\tau}(\text{adt}: T): \mathbb{N} = \text{adt} \text{ match } \{ \\ \text{case cons}_1(v_{1,1}, \dots, v_{1,m_1}) \Rightarrow 1 + \sum_j \text{size}_{\tau_{1,j}}(v_{1,j}) \\ \dots \\ \text{case cons}_n(v_{n,1}, \dots, v_{n,m_n}) \Rightarrow 1 + \sum_j \text{size}_{\tau_{n,j}}(v_{n,j}) \}$$

where $\text{size}_{\tau_2}(x) = 0$ for non-ADT type T_2 . Note that by adding type parameters to the size_{τ} definitions, we can support generic ADT definitions with a finite number of function definitions. It is trivial to (manually) show that each size_{τ} function is terminating. Given expression E , we write $\text{size}_{E.type}(E)$ as $\text{size}(E)$ for brevity. Now given $\mathbf{g}(E_1, \dots, E_n) \in C_{\mathbf{f}}$, consider the expression

$$\text{DEC}(\mathbf{g}(E_1, \dots, E_n)) ::= \text{if}(\text{path}(\mathbf{g}(E_1, \dots, E_n))) \{ \\ \sum_i \text{size}(E_i) < \sum_{v_k \in \mathbf{f}.args} \text{size}(v_k) \\ \} \text{ else } \{ \text{true} \}$$

Let us assume that for all $\mathbf{f} \in P$, $E \in C_{\mathbf{f}}$ and inputs m we have $m[\text{DEC}(E)] \xrightarrow{*} \text{true}$. Now consider the sequence $\text{APP}_0, \text{APP}_1, \dots$ where $\text{APP}_0 ::= \mathbf{f}_0(m[\mathbf{f}_0.args])$, $\text{APP}_i ::= \mathbf{f}_i(A_{i,1}, \dots, A_{i,n_i})$ and $(\text{APP}_i, \text{APP}_{i+1}) \in R$. Given $\sum_j \text{size}(A_{i,j}) \xrightarrow{*} v_i$, we can consider the sequence $v_1 > v_2 > \dots \geq 0$. As no infinite strictly decreasing sequence can exist in the domain of a well-founded relation, Proposition 9 ensures that there is no trace of $m[\mathbf{f}_0.body]$ in which an infinite number of named function application reduction steps occur. Furthermore, we also know that the maximal recursion depth for any named function is upper bounded by h computed as

$$\sum_{v_k \in \mathbf{f}_0.args} \text{size}(m[v_k]) \xrightarrow{*} h.$$

Let us now consider the transformation of P such that all named functions are transformed into anonymous functions and recursive calls are inlined up to depth $h + 1$ (an arbitrary value can be chosen for the result of the deepest call). Clearly, evaluation of the transformed program must achieve the same result as evaluation of P for inputs m . As a result, we obtain program without recursion. Thanks to our assumptions that recursion does not involve recursive type constructors, we can encode data types as well into types of polymorphic lambda calculus. Therefore, the execution of the program reduces to a call-by-value evaluation in polymorphic lambda calculus [31], [28, Chapter 23], in which all evaluations terminate. Consequently, our check ensures termination of the execution. We obtain the following proposition.

Proposition 10. *If there exists no infinite sequence $\text{APP}_0, \text{APP}_1, \dots$ such that $\text{APP}_0 ::= \mathbf{f}_0(m[\mathbf{f}_0.args])$ and $(\text{APP}_i, \text{APP}_{i+1}) \in R$, and all types in P are well-founded, then P terminates on inputs m .*

Verification-based termination. We discussed in previous sections various procedures for showing *validity* of PureScala programs, namely the absence of models such that the program evaluates to a certain value. We can thus employ the $U(\cdot)$ procedure described previously to construct termination proofs

Theorem 11. *If for each $E \in \bigcup_{\mathbf{f} \in P} C_{\mathbf{f}}$ there exists some ϕ_i from $U(\text{DEC}(E))$ such that $\phi_i \in \text{Unsat}$, then P terminates for all inputs.*

Remember that soundness of the $U_{\text{POST}}(\cdot)$ procedure for inductive assumptions requires termination for all inputs of the POST expression. We therefore cannot use this procedure when construction

termination proofs as long as we haven't established termination of POST. Further recall that when considering multiple inductive assumptions, the termination requirement becomes much stronger. Therefore, when considering termination of program P with a set of inductive assumptions $\text{POST}_1, \dots, \text{POST}_m$, we extend the definition of $C_{\mathbf{f}}$ to also contain the set of calls stemming from each relevant POST_i . Formally, given $\mathbf{f}_i(v_{i,1}, \dots, v_{i,n_i})$ associated to POST_i , we extend $C_{\mathbf{f}_i}$ such that

$$\frac{\theta_{\text{args}} = \{v_{i,k} \mapsto u_k \mid u_k \in \mathbf{f}_i.\text{args}\} \quad \theta_{\text{inline}} = \{\mathbf{f}_i(v_{i,1}, \dots, v_{i,n_i}) \mapsto \mathbf{f}_i.\text{body}\}}{C((\theta_{\text{args}} \cup \theta_{\text{inline}})[\text{POST}_i]) \subseteq C_{\mathbf{f}_i}}$$

Note that this extension to $C_{\mathbf{f}_i}$ potentially introduces new pairs into R and relates it as follows with *post-termination*

Proposition 12. *If $\text{POST}_1, \dots, \text{POST}_m$ is not post-terminating for inputs m , then there exist \mathbf{f}_i associated to POST_i , and an infinite sequence $\text{APP}_0, \text{APP}_1, \dots$ such that $\text{APP}_0 ::= \mathbf{f}_i(m[\mathbf{f}_i.\text{args}])$ and $(\text{APP}_i, \text{APP}_{i+1}) \in R$.*

If we can show that for all $\mathbf{f} \in P, E \in C_{\mathbf{f}}$ (using the new definition), and inputs m we have $m[\text{DEC}(E)] \xrightarrow{*} \text{true}$, then P must post-terminate on all inputs (no infinite decreasing and lower-bounded chain can exist). However, $U_{\text{POST}}(\cdot)$ cannot be used to verify the validity of $\text{DEC}(E)$!

Let us consider the call-graph G_p such that $(\mathbf{f}, \mathbf{g}) \in G_p$ iff $\mathbf{f} \in P$ and $\mathbf{g}(E_1, \dots, E_n) \in C_{\mathbf{f}}$. We let G_p^* be the transitive closure of G_p . Given $\mathbf{f} \in P$, we let $\text{Calls}(\mathbf{f}) = \{\mathbf{g} \mid (\mathbf{f}, \mathbf{g}) \in G_p^*\}$ and $\text{Posts}(\mathbf{f}) = \{\text{POST}_i \mid (\mathbf{f}, \mathbf{f}_i) \in G_p^*\}$. If we can show that for all $\mathbf{g} \in \text{Calls}(\mathbf{f}), E \in C_{\mathbf{g}}$, and inputs m we have $m[\text{DEC}(E)] \xrightarrow{*} \text{true}$, then $\text{Posts}(\mathbf{f})$ is post-terminating for all inputs. Indeed, if $(\mathbf{f}, \mathbf{h}) \notin G_p^*$, then \mathbf{h} is entirely irrelevant to post-termination of $\text{Posts}(\mathbf{f})$. These considerations lead to the following statement

Theorem 13. *For $\mathbf{f} \in P$, if $\text{Posts}(\mathbf{f})$ is post-terminating for all inputs, $\text{Posts}(\mathbf{f})$ are valid, and for each $E \in \bigcup_{\mathbf{g} \in P} C_{\mathbf{g}}$ there exists some ϕ_i from $U_{\text{POST}}(\text{Posts}(\mathbf{f}), \text{DEC}(E))$ such that $\phi_i \in \text{Unsat}$, then P terminates for all inputs.*

We can thus construct a termination checking procedure that leverages (independently verified) inductive assumptions in a sound way.

Non-termination. If termination proving fails, we consider the relation $R_{\mathbf{f}} \subseteq R$ such that $(\mathbf{f}(A_1, \dots, A_n), \mathbf{g}(B_1, \dots, B_m)) \in R_{\mathbf{f}}$ iff there exists $\mathbf{g}(E_1, \dots, E_m) \subseteq \mathbf{f}.\text{body}$ such that i) there exists no $((x_1, \dots, x_n) \Rightarrow R) \subseteq \mathbf{f}.\text{body}$ for which $\mathbf{g}(E_1, \dots, E_m) \subseteq R$, and ii) given $\theta = \{v_k \mapsto A_k \mid v_k \in \mathbf{f}.\text{args}\}$ we have $\theta[\text{path}(\mathbf{g}(E_1, \dots, E_m))] \xrightarrow{*} \text{true}$ and $\theta[E_k] \xrightarrow{*} B_k, 1 \leq k \leq m$. Intuitively, for all elements of $R_{\mathbf{f}}$, if an evaluation trace contains the first element of the pair, it must contain the second as well. If we can find a sequence $\text{APP}_1, \dots, \text{APP}_m$ such that $(\text{APP}_i, \text{APP}_{i+1}) \in R_{\mathbf{f}}$ and $\text{APP}_1 = \text{APP}_m$, then for each APP_i , we have a set of arguments that cause non-terminating of the corresponding function. Given some call-graph cycle, the above conditions can be encoded as a satisfiability query and handled by our verification procedure (see below).

6. Counterexamples and Function Equivalence

We have discussed how our procedure can be used to soundly verify properties of PureScala programs, however, its real strength lies in reporting counterexamples to these properties when they exist. We show how our unfolding-based approach can be extended to provide incremental encodings of program over-approximations which correspond to concrete executions when satisfied.

Inputs and equality. Given $(\phi_i, A_i, F_i) \in U(P)$ with $M \models \phi_i$, let us consider the conversion from M to a set of inputs m . Given a

value term $t \in M$, we therefore need a means of translating it back to a PureScala expression. As the sort conversion function \mathcal{S} is not bijective, we further provide the expected type of t . Additionally, as the interpretation of the function corresponding to a $\sigma_{\mathbf{f}}$ -sorted term is given by the interpretation of the associated dispatch symbol, our translation depends on the model M . We inductively define the reverse transformation \mathcal{C}_M as follows

$$\mathcal{C}_M(\text{true}, \text{Boolean}) = \text{true} \quad \mathcal{C}_M(\text{false}, \text{Boolean}) = \text{false}$$

$$\frac{E_1 = \mathcal{C}_M(e_1, T_{i,1}) \quad \dots \quad E_m = \mathcal{C}_M(e_m, T_{i,m})}{\mathcal{C}_M(\text{Cons}_i(e_1, \dots, e_m), T) = \text{cons}_i(E_1, \dots, E_n)}$$

It remains to consider the translation when t has sort $\sigma_{\mathbf{f}}$ (and the expected type is therefore $T = (T_1, \dots, T_n) \Rightarrow T_{res}$. We rely here on the set A_i of function applications witnessed up to this point, and consider the set $A_i(t) \subseteq A_i$ such that for $\text{app} \in A_i(t)$, we have $M \models (\text{app})_b, M[\!(\text{app})_c] = t$, and $(\text{app})_{tpe} = T$. We also assume that we can construct some arbitrary value V_{res} of type T_{res} (note that otherwise, no terminating function exists for T). Given $A_i(t) = \{\text{app}_1, \dots, \text{app}_m\}$ where each $\text{app}_j = (b_j, c_j, tpe, a_j^n)$, we define

$$\frac{E_{j,k} = \mathcal{C}_M(M[\![a_{j,k}]\!], T_k), 1 \leq j \leq m, 1 \leq k \leq n}{R_j = \mathcal{C}_M(M[\![\text{dispatch}_{tpe}(c_j, a_{j,1}, \dots, a_{j,n})]\!], T_{res}), 1 \leq j \leq m}$$

$$\frac{\mathcal{C}_M(t, T) = (v_1 : T_1, \dots, v_n : T_n) \Rightarrow \text{if } (v_1 == E_{1,1} \ \&\& \ \dots \ \&\& \ v_n == E_{1,n}) R_1}{\vdots}$$

$$\text{else if } (v_1 == E_{m,1} \ \&\& \ \dots \ \&\& \ v_n == E_{m,n}) R_m \quad \text{else} \quad V_{res}$$

It is often interesting given a variable v to consider $\mathcal{C}_M(\mathcal{V}(v), v.tpe)$ which we write $\mathcal{C}_M(v)$ for short. We also write $\mathcal{C}_M(\text{vars}(P))$ to describe the inputs $\{v_k \mapsto \mathcal{C}_M(v_k) \mid v_k \in \text{vars}(P)\}$.

Note that the above definition of function-typed term extraction imposes a strong constraint on the kind of first-class functions considered as inputs to the program. However, by disallowing the equality predicate between function-typed expressions in our programs, the above becomes a sound representation of (terminating) functions. Not that our procedure doesn't consider non-terminating functions in program inputs and will not be able to construct V_{res} when no terminating input exists. However, we consider such cases as pathological and these can be identified (and reported) based on the program type definitions.

One can define the $\&\&$ operator through an *if*-expression. Furthermore, meaningful equality predicates exist both in the theory of boolean algebra (\iff) and that of ADTs. However, although we do ensure the \simeq predicate is defined on terms of sort $\sigma_{\mathbf{f}}$, our encoding of first-class functions into fresh constants doesn't ensure a sound congruence relation. Furthermore, equality between functions is a problematic notion in itself. Indeed, full functional equality (in the HOL sense) is undecidable and would therefore make no sense in our operational semantics. Conversely, reference equality would break purity and would require explicit modelling of the heap.

As mentioned previously, syntactic equality of functions can lead to sound models given the existence of inputs. Such a definition of equality can furthermore be efficiently implemented through techniques such as hash-consing. Finally, one can define a corresponding encoding into the formula domain. We start by inductively defining the notion of *simple* PureScala expressions as

$$\text{simple} ::= v \mid \text{true} \mid \text{false} \mid \text{cons}_i(\text{simple}, \dots, \text{simple}) \mid \text{function } \mathbf{f}$$

Note that given a *simple* expression E and a mapping θ from free variables in E to PureScala values, $\theta[E]$ is a value. Now consider the set F obtained from some number of applications

of the FUN rule defined in \rightsquigarrow . Given $fun_1, fun_2 \in F$ where $fun_j = (b_j, f_j, \mathbf{f}_j, c_j^n), j \in \{1, 2\}$, consider the closure substitution $\theta_j = \{\mathcal{V}(\mathbf{f}_j.closures) \mapsto c_j^n\}$. We then inductively define the $eq(E_1, E_2)$ predicate as follows

1. if for $j \in \{1, 2\}$, E_j is *simple* and $\text{vars}(E_j) \subseteq \mathbf{f}_j.closures$,
$$eq(E_1, E_2) \iff \theta_1[C_t(E_1)] \simeq \theta_2[C_t(E_2)];$$
2. else if for $j \in \{1, 2\}$, $E_j = v_j$, then $eq(E_1, E_2) \iff v_1 = v_2$;
3. else if E_1 and E_2 are produced by the same non-terminal in the PureScala grammar, then given their respective children $E_{1,1}, \dots, E_{1,n}$ and $E_{2,1}, \dots, E_{2,n}$,
$$eq(E_1, E_2) \iff eq(E_{1,1}, E_{2,1}) \wedge \dots \wedge eq(E_{1,n}, E_{2,n});$$
4. otherwise, $\neg eq(E_1, E_2)$ must hold.

We define the \sim predicate such that $(f_1, \mathbf{f}_1, c_1^n) \sim (f_2, \mathbf{f}_2, c_2^n)$ corresponds to the clause given by $f_1 \simeq f_2$ iff $eq(\mathbf{f}_1, \mathbf{f}_2)$. We further extend the definition to $fun_1 \sim fun_2$ (b_1 and b_2 are irrelevant). We can now state the following proposition that follows from the freshness of f_1 and f_2

Proposition 14. *For $fun_1, fun_2 \in F$, there must exist some model M such that $M \models fun_1 \sim fun_2$.*

Note that in case No. 1 of the definition of eq , we refer to the associated term $C_t(E_j)$ where E_j has not yet been transformed. However, as E_j does not depend on $\mathbf{f}_j.args$, θ_j is equivalent with $\theta_{\mathbf{f}}$ defined in $\mathcal{I}(app, fun)$ for such expressions, and these can be computed ahead of time. Let us now consider the transformation $(b, E_j) \rightsquigarrow (e_j, \phi_j, A_j, F_j)$ for some given b . Our definition of *simple* expressions ensures that e_j doesn't depend on b , and both ϕ_j and A_j are empty. Further note that the \sim predicate between members of $\theta_j[F_j]$ doesn't depend on b and its choice is therefore irrelevant in our use cases.

Given $x, y \in F$, let $F_{eq}(x, y)$ be the union of the $\theta_j[F_j]$ sets corresponding to each E_j encountered during computation of $x \sim y$ along with x and y themselves. Further let $F_{eq}^*(x, y)$ be the fixpoint where

1. $F_{eq}(x, y) \subseteq F_{eq}^*(x, y)$, and
2. given $x', y' \in F_{eq}^*(x, y)$, $F_{eq}(x', y') \subseteq F_{eq}^*(x, y)$.

One can see that F_{eq}^* must be finite as there can only exist finitely many lambdas within a program. Given a set F , we can thus define the clause set

$$EQ_{\mathbf{f}}(F) = \{x' \sim y' \mid x, y \in F \wedge x', y' \in F_{eq}^*(x, y)\}$$

and extend Proposition 14 to the set F , again ensured by freshness

Proposition 15. *There must exist some model $M \models EQ_{\mathbf{f}}(F)$.*

Lemma 16. *For $fun_1, fun_2 \in F$ where $fun_j = (_, f_j, \mathbf{f}_j, _)$ and \mathbf{f}_j is ground for $j \in \{1, 2\}$, we have $\mathbf{f}_1 = \mathbf{f}_2$ iff some model $M \models f_1 \simeq f_2 \wedge EQ_{\mathbf{f}}(F)$ exists.*

Enumerating relevant functions. We have until now only discussed how a sound congruence relation can be enforced on the tuples fun_1, fun_2 produced by \rightsquigarrow , yet no such tuples exist in F for first-class functions that are *free* in the program. As our programs cannot refer to function equality, it is conservative to consider that the sets of input functions and functions defined within the program do not intersect. Indeed, purity cannot be broken by this constraint. Furthermore, since our definition of function equality depends on the syntax of the function, function-typed input extraction C_M can be easily extended to ensure that if $M[t_1] \neq M[t_2]$ for two $\sigma_{\mathbf{f}}$ sorted terms t_1, t_2 , then $C_M(t_1, T) \neq C_M(t_2, T)$ by adding redundant **if**-expressions to the resulting first-class function bodies. It is

therefore safe to assert that $f_1 \not\approx f_2$ if at least one of f_1, f_2 is free in the program. This observation enables us to avoid explicit modelling of the value associated to each free function in the program.

It remains to discuss how to obtain an exhaustive enumeration of the free functions in P . We introduce here the notion of *sort enumeration*. Given a sort σ , its enumeration consists of the sequence t_0, t_1, \dots of terms with sort σ such that the following conditions are satisfied:

Completeness: for term t of sort σ , there exists $j \in \mathbb{N}$ and some model M such that $M \models t \simeq t_j$.

$\sigma_{\mathbf{f}}$ -exhaustiveness: for any value term t of sort σ with subterm $s \subset t$ of sort $\sigma_{\mathbf{f}}$, there exists $j \in \mathbb{N}$, $s_j \subset t_j$ and some model M such that $M \models t \simeq t_j$ and $M \models s \simeq s_j$.

Unicity: for t_i, t_j in the enumeration such that $i \neq j$ and for model M , we must have $M \models t_i \not\approx t_j$.

It is clear that an enumeration of all value terms of sort σ will satisfy the previous conditions. The proposed definition is more general to enable useful optimizations (consider for example the constant i of integer sort which is a valid enumeration of integers).

Let us now consider the formula ϕ with $v \in \text{vars}(\phi)$ where v has sort σ and t_0, t_1, \dots is an enumeration of σ . Consider the set of terms $E_i(v) = \{s \mid s \subseteq t_j, j \leq i \wedge s \text{ has sort } \sigma_{\mathbf{f}}\}$ and the clause $\eta_i(v) = v \in \{t_j \mid j \leq i\}$. The completeness condition of sort enumerators ensures satisfiability is preserved and gives us the following proposition:

Proposition 17. *If $M \models \phi$ exists, then there exists $j \in \mathbb{N}$ and model M_j such that $M_j \models \phi \cup \eta_j(v)$.*

More importantly, sort enumerators give us the means to consider the set of all relevant functions, leading to the following lemma:

Lemma 18. *Given a model $M \models \eta_i(v)$, there is no term $s \subseteq M[v]$ with sort $\sigma_{\mathbf{f}}$ such that $M \models s \notin E_i(v)$.*

When considering multiple sort enumerations jointly, we require that the sets of constants introduced by each enumeration be disjoint. This property can be trivially ensured through freshening.

Consider the set A obtained through some number of applications of the APP rule defined in \rightsquigarrow , and V the set of constants obtained by transforming the free variables in the program such that constants in A, F and $EQ_{\mathbf{f}}(F)$ either belong to V or were introduced by the \rightsquigarrow transformation. As the sort $\sigma_{\mathbf{f}}$ is not truly related to a theory of first-class functions, we must consider function terms contained *within* free functions. Indeed, free function application results that contain terms with sort $\sigma_{\mathbf{f}}$ are not considered during our enumeration and therefore won't belong to any $E_i(v)$. Now consider $app \in A$ where $app = (b, c, tpe, a^n)$: if the clause $c \in E_i(v)$ holds, then we must enumerate the sort of $\text{dispatch}_{tpe}(c, a_1, \dots, a_n)$ (noted $\text{disp}(app)$ hereafter) as it may contain further relevant functions. We can therefore define the set of all potentially relevant function terms as

$$E_{all,i}(V, A) = \bigcup_{v \in V} E_i(v) \cup \bigcup_{app \in A} E_i(\text{disp}(app)).$$

Remember that we want to enforce distinctness with respect to the function tuples in F , thus leading to the definition

$$EQ_{f,i}(V, A, F) = \{f \not\approx (fun)_f \mid f \in E_{all,i}(V, A), fun \in F\}.$$

Proposition 19. *There must exist model $M \models EQ_{f,i}(V, A, F)$.*

We have discussed how $\eta_i(v)$ will bind the terms in $E_i(v)$ to v , however we cannot assert $\eta_i(\text{disp}(app))$ directly as $\text{disp}(app)$ is only free when $(app)_c \in E_{all,i}(V, A)$ holds. We therefore want a *conditional* binding for $E_i(\text{disp}(app))$, given by the clause set

$\eta_{A,i}(V, A)$ defined as follows

conditionally_i(app, S) = (app_c) $\in S \implies \eta_i(\text{disp}(app))$

$\eta_{A,i}(V, A) = \{ \text{conditionally}_i(app, E_{all,i}(V, A)) \mid app \in A \}$.

Finally, given $(\phi_i, A_i, F_i) \in U(\mathbb{P})$ and $V = \{ \mathcal{V}(v) \mid v \in \text{vars}(\mathbb{P}) \}$, we define the following clause set

$\psi_i = EQ_x(F_i) \cup EQ_{f,i}(V, A_i, F_i) \cup \eta_i(V) \cup \eta_{A,i}(V, A_i)$.

Over-approximating the program. Given the uninterpreted nature of the various dispatch_{*tpc*} function symbols, we want to consider an over-approximation of the program where only paths for which all relevant inlinings have taken place are allowed. We define $D_i = \{ (app_j, fun_j) \mid j \leq i \}$, the set of all application-function pairs that have been considered (*i.e.* dispatched) up to this point. For $app \in A_i$, we define the set of functions that have been considered as dispatch targets for app

$T_i(app) = \{ (fun_j)_f \mid (app_j, fun_j) \in D_i, app_j = app \}$.

Now consider the case where the caller $(app)_c$ is in fact *free* in the program (*i.e.* comes from a free variable). No dispatch for app is necessary as the interpretation for $(app)_c$ defined by the reverse transformation \mathcal{C}_M corresponds to the semantics given by dispatch. We can therefore define the set of handled callers for app

$T_{all,i}(app) = T_i(app) \cup E_{all,i}$.

Remember that each clause in ϕ_i is of the shape $b \implies c$ for some constant b . If $(app)_c \notin T_{all,i}(app)$ holds, then all parts of the program that depend on the result of the function application associated to app must be disallowed as the relevant function inlining hasn't taken place yet. This condition extends to each $app \in A_i$ and leads to the clause set

$\rho_i = \{ (app)_c \notin T_{all,i}(app) \implies \neg(app)_b \mid app \in A_i \}$.

We assume a fair selection strategy of (app_i, fun_i) at each inlining step, so for any $app \in A_i, fun \in F_i$ there exists $j \in \mathbb{N}$ such that $app = app_j$ and $fun = fun_j$. We can thus state the following:

Theorem 20. *Given a model $M \models \phi_i \cup \psi_i \cup \rho_i$, we have $\mathcal{C}_M(\text{vars}(\mathbb{P}))[\mathbb{P}] \xrightarrow{*} \text{false}$.*

Theorem 21. *If inputs m exist such that $m[\mathbb{P}] \xrightarrow{*} \text{false}$, then there exists $i \in \mathbb{N}$ and model M such that $M \models \phi_i \cup \psi_i \cup \rho_i$.*

Corollary 22. *For inductive assumption POST where POST is valid and terminates on all inputs, given $(\phi_i, A_i, F_i) \in U_{\text{POST}}(\mathbb{P})$ with associated ψ_i, ρ_i and model $M \models \phi_i \cup \psi_i \cup \rho_i$, we have $\mathcal{C}_M(\text{vars}(\mathbb{P}))[\mathbb{P}] \xrightarrow{*} \text{false}$.*

Conversely, if inputs m exist such that $m[\mathbb{P}] \xrightarrow{} \text{false}$, there exists $j \in \mathbb{N}$ and model M such that given $(\phi_j, A_j, F_j) \in U_{\text{POST}}(\mathbb{P})$ with associated ψ_j, ρ_j , we have $M \models \phi_j \cup \psi_j \cup \rho_j$.*

7. Algebraic Data Type Invariants

Given an ADT type \mathbb{T} , we define an algebraic data type invariant as a PureScala expression INV with a single free variable $x : \mathbb{T}$ that is universally quantified. Semantically, the ADT invariant implies that for any *valid* instance $E_{\mathbb{T}}$ of type \mathbb{T} , we must have $\{ x \mapsto E_{\mathbb{T}} \} \llbracket \text{INV} \rrbracket \xrightarrow{*} \text{true}$. We describe how the concept of sort enumeration can be further leveraged to support finding models that agree with ADT invariants.

It is interesting to note that ADT invariants do not increase the expressivity of contracts for first-order functions (though they do decrease the annotation burden). However, they enable some level of specification for first-class functions by leveraging the type system. There exists a clear relation between our ADT invariants and refinement types. However, as our refinement is *named*, we need only consider \mathbb{T} -typed instance construction points to ensure a valid

$$\begin{array}{c}
 \text{true} \triangleright \text{INV} \qquad \qquad \qquad \text{false} \triangleright \text{INV} \\
 \\
 \frac{E_1 \triangleright \text{INV} \quad \dots \quad E_m \triangleright \text{INV} \quad \text{cons}_i(E_1, \dots, E_m) : \mathbb{T} \quad \{ x \mapsto \text{cons}_i(E_1, \dots, E_m) \} \llbracket \text{INV} \rrbracket^* \longrightarrow \text{true}}{\text{cons}_i(E_1, \dots, E_m) \triangleright \text{INV}} \\
 \\
 \frac{E_1 \triangleright \text{INV} \quad \dots \quad E_m \triangleright \text{INV} \quad \text{cons}_i(E_1, \dots, E_m) : \mathbb{T}_2, \mathbb{T}_2 \neq \mathbb{T}}{\text{cons}_i(E_1, \dots, E_m) \triangleright \text{INV}} \\
 \\
 \frac{\text{f} : (\mathbb{T}_1, \dots, \mathbb{T}_n) \Rightarrow \mathbb{T}_{res} \quad \forall E_1 : \mathbb{T}_1, \dots, E_n : \mathbb{T}_n, \text{f}(E_1, \dots, E_n)^* \longrightarrow v \triangleright \text{INV}}{\text{f} \triangleright \text{INV}}
 \end{array}$$

Figure 5. Inference rules for the $E \triangleright \text{INV}$ relation.

program typing. We therefore generate and prove corresponding verification conditions and show termination for all inputs of the given INV expression, thus intuitively ensuring both progress and preservation in the corresponding refinement type system.

Unlike relevant function enumeration, our notion of ADT invariants is not based on some need to relate terms resulting from the transformation \rightsquigarrow but on the expressions themselves. It is therefore not so clear what the set of relevant instances of \mathbb{T} actually consists of. We clarify the notion through a *holds* relation on PureScala values where $E \triangleright \text{INV}$ is defined in Figure 5. Given a set of inputs m to a program \mathbb{P} returned by \mathcal{C}_M , we say INV holds on m iff for all $v \in \text{vars}(\mathbb{P})$, we have $m \llbracket v \rrbracket \triangleright \text{INV}$.

Let us consider the formula ϕ with $v \in \text{vars}(\phi)$ where v has sort σ . For any term t of sort σ , we are naturally interested in all subterms $s \subseteq t$ of sort $\sigma_{\mathbb{T}}$ where $\sigma_{\mathbb{T}} = \mathcal{S}(\mathbb{T})$. We therefore extend the definition of sort enumeration t_0, t_1, \dots for sort σ by adapting the σ_x -exhaustiveness condition to the $\sigma_{\mathbb{T}}$ sort:

$\sigma_{\mathbb{T}}$ -exhaustiveness: for any value term t of sort σ with subterm $s \subseteq t$ of sort $\sigma_{\mathbb{T}}$, there exists $j \in \mathbb{N}, s_j \subseteq t_j$ and some model M such that $M \models t \simeq t_j$ and $M \models s \simeq s_j$.

We define the set $S_i(v) = \{ s \mid s \subseteq t_j, j \leq i \wedge s \text{ has sort } \sigma_{\mathbb{T}} \}$ by analogy with $E_i(v)$ of all subterms in t_0, \dots, t_i that have sort $\sigma_{\mathbb{T}}$. Note that $\eta_i(v)$ ensures terms in $S_i(v)$ are bound as well as those in $E_i(v)$. Proposition 17 trivially extends to enumerations with $\sigma_{\mathbb{T}}$ -exhaustiveness and we can further state the following lemma:

Lemma 23. *Given a model $M \models \eta_i(v)$, there is no term $s \subseteq M \llbracket v \rrbracket$ with sort $\sigma_{\mathbb{T}}$ such that $M \models s \notin S_i(v)$.*

As with function applications, the definition of ADT invariants leads us to consider $(b_{\mathbb{T}}, \text{INV}) \rightsquigarrow (inv, \phi_{\mathbb{T}}, A_{\mathbb{T}}, F_{\mathbb{T}})$ where $b_{\mathbb{T}}$ is a fresh constant. Given the subterm $s \in S_i(v)$ that comes from enumeration term t_j , we define the substitution $\theta_s = \{ \mathcal{V}(x) \mapsto s \}$. Intuitively, as long as $v \simeq t_j$, the substitutions $\theta_s \llbracket \{ inv \} \cup \phi_{\mathbb{T}} \rrbracket, \theta_s \llbracket A_{\mathbb{T}} \rrbracket$ and $\theta_s \llbracket F_{\mathbb{T}} \rrbracket$ correspond to taking the transformation of $\mathbb{P} \ \&\& \ \{ x \mapsto \mathcal{C}_M(s) \} \llbracket \text{INV} \rrbracket$. This leads to the definition

$$\mathcal{P}_{\text{INV}}(s) = (\theta_s \llbracket \phi_{\mathbb{T}} \rrbracket \cup \{ b_{\mathbb{T}} \iff (v \simeq t_j), \\ b_{\mathbb{T}} \implies \theta_s \llbracket inv \rrbracket \}, \theta_s \llbracket A_{\mathbb{T}} \rrbracket, \theta_s \llbracket F_{\mathbb{T}} \rrbracket).$$

Note that, again, we preserve the $b \implies c$ structure of the resulting clause set. We require that $b_{\mathbb{T}}$ be fresh for each $s \in S_i(v)$ as it is fully defined by the resulting clause set. As in the relevant function terms case, the set of all relevant $\sigma_{\mathbb{T}}$ -sorted terms is given by

$$S_{all,i}(V, A) = \bigcup_{v \in V} S_i(v) \cup \bigcup_{app \in A} S_i(\text{disp}(app)).$$

Recall that our definition of $\mathcal{C}_M(t, (\mathbb{T}_1, \dots, \mathbb{T}_n) \Rightarrow \mathbb{T}_{res})$ supposes the ability to generate a value \mathbb{V}_{res} of type \mathbb{T}_{res} . As our goal is to

produce inputs m such that INV holds on m , we must ensure that $V_{res} \triangleright INV$ for all first-class functions in $\mathcal{C}_M(\text{vars}(\mathbb{P}))$. The proof of Theorem 20 shows that we can construct such a value by extending u_0 with $(\{b_{\top}, inv\} \cup \phi_{\top}, A_{\top}, F_{\top})$. However, if $\neg INV$ is valid, namely no value E_{\top} exists such that $\{x \mapsto E_{\top}\} \llbracket INV \rrbracket \xrightarrow{*} \mathbf{true}$, then ϕ_i may be Unsat even though no value of type T is required in m . Given $v \in \text{vars}(\mathbb{P})$, one can easily determine whether $v.tpe$ will require an instance V_{\top} of type T inside of $\mathcal{C}_M(\text{vars}(\mathbb{P}))$. We say $\text{vars}(\mathbb{P})$ depend on T iff there exists a $v \in \text{vars}(\mathbb{P})$ such that $v.tpe$ requires an instance V_{\top} .

Based on these considerations, we extend the $U(\mathbb{P})$ procedure to $U_{INV}(\mathbb{P}) = u_0, u_1, \dots$ in the following way. Given $u'_0 \in U(\mathbb{P})$ and $(b_p, INV) \rightsquigarrow (inv, \phi_p, A_p, F_p)$, if (and only if) $\text{vars}(\mathbb{P})$ depend on T , we let $u_0 = u'_0 \cup (\{inv\} \cup \phi_p, A_p, F_p)$ and add x to $\text{vars}(\mathbb{P})$. Furthermore, we extend $\mathcal{C}_M(t, (T_1, \dots, T_n) \Rightarrow T_{res})$ such that $\mathcal{C}_M(x)$ is used when constructing V_{res} if T_{res} depends on T . Otherwise, we have $u_0 = u'_0$ and $\text{vars}(\mathbb{P})$ are left unchanged. Finally, u_i is defined given u_{i-1} as

$$u_i = u_{i-1} \cup \mathcal{I}(app_i, fun_i) \cup \bigcup_{s \in S_{all,i}(\text{vars}(\mathbb{P}), A_i)} \mathcal{P}_{INV}(s).$$

Clearly, given $(\phi_i, A_i, F_i) \in U_{INV}(\mathbb{P})$, the extension to $\phi_i \cup \psi_i \cup \rho_i$ only adds constraints to potential models, giving us

Proposition 24. For $(\phi_i, A_i, F_i) \in U_{INV}(\mathbb{P})$ with associated ψ_i, ρ_i , given $M \models \phi_i \cup \rho_i \cup \psi_i$, then $\mathcal{C}_M(\text{vars}(\mathbb{P})) \llbracket \mathbb{P} \rrbracket \xrightarrow{*} \mathbf{false}$.

The more interesting aspect of the updated procedure resides in the kind of satisfying models that are produced. Indeed, $\mathcal{P}_{INV}(s)$ for $s \in S_{p,i}$ ensures that INV is considered for all relevant ADT terms, thus leading to the following statements:

Theorem 25. For $(\phi_i, A_i, F_i) \in U_{INV}(\mathbb{P})$ with associated ψ_i, ρ_i , if $M \models \phi_i \cup \psi_i \cup \rho_i$, then INV holds on $\mathcal{C}_M(\text{vars}(\mathbb{P}))$.

Theorem 26. If inputs m exist such that INV holds on m and $m \llbracket \mathbb{P} \rrbracket \xrightarrow{*} \mathbf{false}$, then there exists $(\phi_i, A_i, F_i) \in U_{INV}(\mathbb{P})$ with associated ψ_i, ρ_i such that $M \models \phi_i \cup \psi_i \cup \rho_i$ exists.

It is clear that the above results can be extended to an arbitrary set INV_1, \dots, INV_m of ADT invariants as all \mathcal{P}_{INV_i} are independent. Furthermore, one can safely use U_{POST} in conjunction with U_{INV} as both procedures are independent. Note that INV_1, \dots, INV_m cannot introduce relevant non-termination as unsoundness will only arise if $POST_1, \dots, POST_n$ are not post-terminating for all inputs.

8. Explicit First-Order Quantification

To specify assumptions and invariants about first-class functions, we typically wish to describe the behavior of these functions on infinite families of inputs. For this purpose, we introduce explicit universal first-order quantifiers in our specification language. Instead of relying on SMT solvers to instantiate quantifiers, our system performs quantifier instantiation internally. Thanks to this approach, our system can report counterexamples for some of the queries on which our SMT solvers return “Unknown” in their default configuration. While we acknowledge that the use of quantifiers is expensive, it is a price to pay for modular reasoning about first-class functions. Note that our users need not rely on explicit quantifiers as often as in some other verification approaches: constructs such as recursion and ADT invariants can be thought as particularly structured forms of quantification, which our system handles more efficiently, as discussed in previous sections.

Consider the PureScala expression $\forall x_1, \dots, x_m. \text{PROP}$. Clearly, no reasonable evaluation semantics can be attached to \forall , so we define its semantics as follows: given inputs m , we say $\forall x_1, \dots, x_m. \text{PROP}$ holds on m iff for every mapping

$$m_{\text{QUANT}} = \{x_k \mapsto E_k \mid E_k : x_k.tpe, E_k \text{ ground}, 1 \leq k \leq m\},$$

we have $(m \cup m_{\text{QUANT}}) \llbracket \text{PROP} \rrbracket \xrightarrow{*} \mathbf{true}$. Now assume that x_1 to x_m only appear in argument position to function calls in PROP and $x_k.tpe$ is first-order for $1 \leq k \leq m$ (i.e. no instance of $x_k.tpe$ contains a function). Let $(b_p, \text{PROP}) \rightsquigarrow (p, \phi_p, A_p, F_p)$. Given function type $T = (T_1, \dots, T_n) \Rightarrow T_{res}$, consider the set $G_{T,j}(A_p)$, $1 \leq j \leq n$ of tuples (b_G, t) where b_G corresponds to the condition under which t is a valid quantifier-free instantiation for argument position j . Further consider the set $X_k(A_p)$, $1 \leq k \leq m$ of tuples (b_X, s) where b_X is the condition under which s is a valid quantifier-free instantiation for the quantified variable x_k . Given $(b_k, s_k) \in X_k(A_p)$, $1 \leq k \leq m$, we let

$$\text{inst}((b_1, s_1), \dots, (b_m, s_m)) = (b_1 \wedge \dots \wedge b_m, \{ \mathcal{V}(x_k) \mapsto s_k \mid 1 \leq k \leq m \}).$$

For each $app \in A_p$ where $app = (b_{\top}, T, a^n)$ and argument position $1 \leq j \leq n$, consider the following set constraints

$$\text{INIT-G} \frac{a_j \text{ quantifier-free}}{(b, a_j) \in G_{T,j}(A_p)} \quad \text{INIT-X} \frac{a_j = \mathcal{V}(x_k)}{G_{T,j}(A_p) = X_k(A_p)}$$

$$\text{EXPAND} \frac{\begin{array}{c} a_j \text{ not quantifier-free} \\ (b_1, s_1) \in X_1(A_p) \quad \dots \quad (b_m, s_m) \in X_m(A_p) \\ (b_s, \theta_s) = \text{inst}((b_1, s_1), \dots, (b_m, s_m)) \end{array}}{(b \wedge b_s, \theta_s \llbracket a_j \rrbracket) \in G_{T,j}(A_p)}$$

One should note here that the sets $G_{T,j}(A_p)$ and $X_k(A_p)$ are not necessarily finite. We let $\mathcal{X}(A_p) = X_1(A_p) \times \dots \times X_m(A_p)$ and given $x \in \mathcal{X}(A_p)$ where $x = ((b_1, s_1), \dots, (b_m, s_m))$, we write $\text{inst}(x)$ for $\text{inst}((b_1, s_1), \dots, (b_m, s_m))$.

Let us now define the sequence $Q(A_p) = q_0, q_1, \dots$ where $q_i = (\mathcal{X}_i, A_i)$ such that both \mathcal{X}_i and A_i are finite. We define $\mathcal{X}_{\text{INIT}}(A_p)$ as the set obtained using only the inference rules INIT-G and INIT-X on A_p . Note that $\mathcal{X}_{\text{INIT}}(A_p)$ is guaranteed finite. We therefore let $q_0 = (\mathcal{X}_{\text{INIT}}(A_p), A_p)$. Let us now consider $x_i \in \mathcal{X}_{i-1}$ where $(b_s, \theta_s) = \text{inst}(x_i)$. We define A_i as

$$A_i = A_{i-1} \cup \{ (b \wedge b_s, \theta_s \llbracket c \rrbracket, tpe, \theta_s \llbracket a^n \rrbracket) \mid (b, c, tpe, a^n) \in A_{i-1}, a^n \text{ not quantifier-free} \}$$

and let $\mathcal{X}_i = \mathcal{X}_{\text{INIT}}(A_i)$. Note that this definition of \mathcal{X}_i corresponds to applying a single time the EXPAND rule to the set \mathcal{X}_{i-1} and computing the (finite) fixpoint with respect to INIT-G and INIT-X.

Lemma 27. For $Q(A_p) = (\mathcal{X}_0, _), (\mathcal{X}_1, _), \dots, \bigcup_i \mathcal{X}_i = \mathcal{X}(A_p)$.

Note that given $(\mathcal{X}_i, A_i) \in Q(A_p)$, $x \in \mathcal{X}_i$, $(b_k, s_k) \in x$ and some subterm $t = \text{dispatch}_{tpe}(\dots) \subseteq s_k$, there must always exist a corresponding $app \in A_i$ such that $t = \text{disp}(app)$ and $\models b_k \iff (app)_b$ as $\mathcal{X}_{\text{INIT}}$ does not introduce new terms. This ensures that the different underlying $U(\cdot)$ procedures can simply consider the set A_i for their various use cases as it faithfully represents the set of dispatches taking place in the formulas.

Given $x \in \mathcal{X}_i$ and fresh constant b_q , consider the transformation $(b_q, \text{PROP}) \rightsquigarrow (q, \phi_q, A_q, F_q)$ and $(b_s, \theta_s) = \text{inst}(x)$. The substituted sets $\theta_s \llbracket \{q\} \cup \phi_q \rrbracket$, $\theta_s \llbracket A_q \rrbracket$ and $\theta_s \llbracket F_q \rrbracket$ correspond to the instantiation of PROP with the quantifier-free terms s_1 to s_m . We want this instantiation to be conditioned on b_1 to b_m to ensure the quantifier-free terms are well-formed. Based on this, we define

$$\mathcal{P}_{\text{QUANT}}(x) = (\theta_s \llbracket \phi_q \rrbracket) \cup \{ b_q \iff bs, b_q \implies \theta_s \llbracket q \rrbracket, \theta_s \llbracket A_q \rrbracket, \theta_s \llbracket F_q \rrbracket \}.$$

We now define the sequence $U_{\text{QUANT}}(\text{PROP}) = u_1, u_2, \dots$ of triplets $u_i = (\phi_i, A_i, F_i)$ where $u_0 = (\{b_q, q\} \cup \phi_q, A_q, F_q)$. Note that this simply corresponds to introducing the quantifier-free instantiations $(true, \mathcal{V}(x_k))$ into each X_k for $1 \leq k \leq m$, a sound instantiation for non-empty sort $\mathcal{S}(x_k.tpe)$. For u_i , we nondeterministically chose one of the following:

1. $u_i = u_{i-1} \cup \mathcal{I}(app_i \in A_{i-1}, fun_i \in F_{i-1})$,
2. $u_i = u_{i-1} \cup \mathcal{P}_{\text{QUANT}}(x_i \in \mathcal{X}_{\text{INIT}}(A_i))$.

We assume a fair alternation between the two presented options and fair selections of app_i, fun_i, x_i . Note that for $(\mathcal{X}_j, A_j) \in Q(A_q)$, there exists $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$ such that $A_j \subseteq A_i$ and $\mathcal{X}_j \subseteq \mathcal{X}_{\text{INIT}}(A_i)$.

Theorem 28. *Given $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$, if $\phi_i \in \text{Unsat}$, then there exist no inputs m on which $\forall x_1, \dots, x_m. \text{PROP}$ holds.*

Let us now discuss the parallel between our definition of $\mathcal{X}(A_p)$ and the work presented in [10]. Consider the clause set ϕ_p where $\mathcal{V}(x_k), 1 \leq k \leq m$ are quantified. As our transformation introduces a single uninterpreted function symbol for each function type \mathbb{T} , our set constraints correspond exactly to those discussed in [10], with the addition of conditionals. Note that we do not consider clause-indexed instantiations for the quantified variables as the notion of *clause* is not so clear in PureScala, yet this simplification only translates into extra (sound) instantiations in our procedure. The conditional associated to each $\mathcal{P}_{\text{QUANT}}$ instantiation corresponds to a *path condition* under which each instantiation is relevant. As long as PROP falls into a fragment that is covered in [10] (namely only pseudo-macro function definitions), our preserves Theorem 1 from [10] as negated *path conditions* will translate into clause satisfiability for any quantifier-free instantiation of a given clause, therefore rendering the associated quantifier-free instantiations irrelevant. Therefore, there are interesting fragments for which U_{QUANT} constitutes a semi-decision procedure.

Remember from our definition of U_{QUANT} that there may exist $app_q \in A_0$ such that app is not quantifier-free. Now consider $u_i = u_{i-1} \cup \mathcal{I}(app_q, fun_i)$ for some $i \in \mathbb{N}$; given our definition of $\mathcal{X}_i = \mathcal{X}_{\text{INIT}}(A_i)$, if $\mathcal{I}(app_q, fun_i)$ introduces an $app \in A_i$ with argument a_j such that $a_j = \mathcal{V}(x_k)$, \mathcal{X}_i will satisfy the new set constraint given by INIT-X . In other words, our instantiation procedure will automatically account for the implicit universal quantification given by the definition of $(app)_\mathbb{T}$, namely

$$\forall x_1, \dots, x_n. \mathbb{f}(x_1, \dots, x_n) = \{ \mathbb{f}.args \mapsto x^n \} [\mathbb{f}.body].$$

Let us now extend the procedure to model finding. Consider $(\mathcal{X}_i, A_i) \in Q(A_p)$ and model M . We say M *finitely interprets* \mathcal{X}_i iff for each $x \in \mathcal{X}_i$ with $(b_s, \theta_s) = \text{inst}(x)$, we either have

1. $M \models \neg b_s$, or
2. for each $app_q \in A_i$ where $app_q = (b_q, c_q, tpe_q, a_q^n)$ such that app_q is *not* quantifier-free, there exists $app_g \in A_i$ where $app_g = (b_g, c_g, tpe_g, a_g^n)$ such that $M \models b_q \iff b_g$, $M \models \theta_s \llbracket c_q \rrbracket \simeq c_g$, $tpe_q = tpe_g$ and $M \models \theta_s \llbracket a_q^n \rrbracket \simeq a_g^n$.

Lemma 29. *For $(\mathcal{X}_i, A_i) \in Q(A_p)$ and model M , if M finitely interprets \mathcal{X}_i , then M finitely interprets \mathcal{X}_{i+1} .*

As \mathcal{X}_i and A_i are finite, the above conditions can be encoded into a clause set γ_i . Note that given a model $M \models \gamma_i$, then for all $j > i$, we have $M \models \gamma_j$. For a given $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$, we already had associated ψ_i, ρ_i , and we can now consider an associated γ_i obtained from A_i and $\mathcal{X}_{\text{INIT}}(A_i)$. These considerations lead to the following theorem

Theorem 30. *Given $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$ with associated ψ_i, ρ_i, γ_i , and $M \models \phi_i \cup \psi_i \cup \rho_i \cup \gamma_i$, if $\forall x_1, \dots, x_m. \text{PROP}$ is in a semi-decidable fragment, then $\forall x_1, \dots, x_m. \text{PROP}$ holds on $\mathcal{C}_M(\text{vars}(\forall x_1, \dots, x_m. \text{PROP}))$.*

One should note at this point that all first-class functions contained within $\mathcal{C}_M(\cdot)$ have finite range. As our exploration of relevant arguments is complete for semi-decidable fragments, if inputs where all first-class functions have finite range exists, then these ranges will eventually be covered by the considered arguments, leading

to the following (conditional) completeness guarantee for model finding

Theorem 31. *If there exist inputs m such that for each first-class function $\mathbb{f} \in m$, the set $\{ \mathbb{v} \mid \mathbb{f}(E_1, \dots, E_n) \xrightarrow{*} \mathbb{v} \}$ is finite and $\forall x_1, \dots, x_m. \text{PROP}$ is in a semi-decidable fragment, then there exists $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$ with associated ψ_i, ρ_i, γ_i , and model M such that $M \models \phi_i \cup \psi_i \cup \rho_i \cup \gamma_i$.*

It is again trivial to extend $U_{\text{QUANT}}(\cdot)$ to multiple $\text{PROP}_1, \dots, \text{PROP}_m$. Note that when considering U_{QUANT} in conjunction with U_{INV} , we must ensure that our assumption on type inhabitation is satisfied, thus conditioning soundness of validity proofs to satisfiability of the relevant INV_i expression.

It is also interesting to consider the procedure that takes a quantifier-free program P as argument and shows that to inputs m exist such that $\text{PROP}_1, \dots, \text{PROP}_m$ hold on m and $m \llbracket P \rrbracket \xrightarrow{*} \text{false}$. Given the transformation $(b_p, P) \rightsquigarrow (p, \phi_p, A_p, F_p)$, it suffices to let $(\phi_p \cup \{b_p, \neg p\}, A_p, F_p) \subseteq u_0$ in $U_{\text{QUANT}}(\{\text{PROP}_1, \dots, \text{PROP}_m\}, P)$ to obtain such a procedure. However, termination considerations now appear, much as in the U_{POST} case. These are handled similarly to U_{POST} by extending the call-graph with new relevant edges and showing that the considered measure remains strictly decreasing.

9. Implementation and Evaluation

We have implemented the techniques described in this paper as a plugin for the scalac compiler that accepts a subset of Scala and performs verification of such programs. We illustrate the behavior of our verifier on a sample of benchmarks to indicate that our techniques are of practical interest; a more detailed description of implementation and evaluation is beyond the scope of this anonymous submission.

Table 1 shows the results we choose to highlight. Most verification benchmarks feature some use of higher-order functions and a fair selection display useful interactions between ADT invariants and first-order quantifiers. Our benchmarks present a non-trivial List library with the usual higher-order operators which feature verified contracts describing various properties such as associativity (map, flatMap), equivalence (exists, forall, folds, etc.), and the monadic laws for flatMap. It is interesting to note that we can produce interface-like features by passing around first-class functions. Consider for example the following signature

```
def mergeSort[T](list: List[T])(key: T => Z): List[T] = { ... }
ensuring { res =>
  isSorted(res)(key) && content(list)(key) == content(res)(key) }
```

where the key function provides a mapping into the well-ordered domain of integers. Multiset contents and sortedness are then defined with respect to the provided key function.

ADT invariants have enabled a significant decrease of the annotation burden in certain involved benchmarks such as binary search tree insertion verification (Tree). First-order quantifiers are featured in the Tree, MergeSort (where they are used to encode higher-order functional equality), ArraySearch and AssocCommut (various mathematical properties of associative and commutative functions) benchmarks. Note that although the performance of these benchmarks is slower than those without quantifiers, such verification tasks remain tractable and feature a very high degree of automation. In terms of performance, our implementation of the quantifier instantiation procedure is in many cases not as fast as implementations that would use internal data structures of SMT solvers, but we have identified cases where quantifier instantiation techniques in the solvers we use (Z3 and CVC4) do not succeed in reporting models for quantified conjectures, whereas our approach does.

Our termination prover shows termination of all listed verification benchmarks, though it requires the introduction of ghost

Table 1. Summary of evaluation results for verification, featuring lines of code, (V)alid, (I)nvaid and (U)nkown verification conditions, as well as running time of the procedure.

Operation	LoC	V	I	U	Time (s)
List	820	242	0	0	6.85
Tree	95	15	0	1	1.10
StateMonad	212	19	0	0	0.50
OptionMonad	49	9	0	0	0.39
ExceptionMonad	61	6	0	0	0.36
FoldAssociative	104	24	0	0	1.40
BalancedBrackets	225	39	0	0	1.12
ArraySearch	166	29	2	0	20.64
MergeSort	67	13	0	0	11.64
AssocCommot	44	3	2	0	13.80
Total	1843	399	4	1	57.80

Table 2. Summary of evaluation results for termination checker, featuring lines of code, (P)roved functions, (N)on-terminating inputs reported and (U)nkown, as well as running time of our tool.

Operation	LoC	P	N	U	Time (s)
NNF	101	14	0	0	13.49
Ackermann	10	2	0	0	1.54
QuickSort	62	10	0	0	8.48
RedBlackTree	60	8	0	0	1.25
LambdaCalculus	53	4	1	0	11.78
CyclicReasoning	30	1	3	0	8.13
HOTransformations	41	3	0	0	5.43
McCarthy91	10	0	0	1	(150.43)
Total	357	42	4	1	50.10

variables to handle binary search and merge sort as the relevant measures are entirely non-trivial. We further list certain interesting termination benchmarks that showcase the flexibility of our tool. The evaluation results for these can be found in Table 2. Our prover can also report non-terminating inputs in certain classic cases of non-termination. For example, our prover is able to show that evaluation is non-terminating in the untyped lambda calculus (Appendix D.1) and produces the non-terminating input $(\lambda x.xx)(\lambda x.xx)$, namely the well-known omega term. Our termination checker defers non-termination checking until all positive techniques have failed to produce a proof, contributing to larger execution times in non-terminating cases. Overall, we obtain a good tradeoff between automation and performance by using a cascading model of termination checking that starts by invoking simple (and fast) procedures and then discharges failed checks to more powerful techniques. In practice, a significant portion (almost half) of our termination proofs are constructed without requiring any solver interaction, justifying this incremental approach.

10. Related Work

Verification of higher-order functional programs has gained recent traction through different significant approaches.

Model checkers for higher-order programs. Dependent type systems with refinement types such as *Liquid Types* [32, 39] and approaches based on model checking higher-order recursion schemas [15, 16, 27] are powerful verification techniques that elegantly

handle higher-order functions. These procedures can benefit from counterexample-guided abstraction refinement (CEGAR) for high levels of automation and even boast relative completeness results for program typing [36]. It can however prove quite challenging to identify the cause of proof failure [38]. Note that our ADT invariants can be viewed as some form of refinement types [13, 37]. Recent work has shown that shape analysis of symbolic execution traces can prove effective at identifying contract violations in a (relatively) complete manner [23, 24]. These techniques, although similar in their aim to ours, use significantly different approaches and are therefore adapted to orthogonal problems such as dynamically typed languages and mutable state. Other techniques focus on unfolding first-order recursive functions [34, 35] or treat the simpler case of functions without datastructures or quantifiers [40].

Theorem provers and solvers. Automatic HOL theorem provers such as LEO II [4] and Satallax [8] feature semi-decidability results for certain HOL semantics. Given some suitable encoding, these systems can automatically verify properties in the fragment we discussed, however, their effectiveness on benchmarks with heavy usage of data types and integers is expected to be limited [33]. Saturation-based first-order theorem provers [1, 25] can efficiently handle *first-order* logic with equality [2, 17] but handling theories and higher-order functions remains a difficulty despite progress [29]. Interactive theorem provers such as Isabelle/HOL [26] and Coq [5] enable large scale real world verification efforts at the cost of automation. Partial automation [7], and counterexample finders for higher-order functions [6, 22] further improve the usability of these tools, though user interaction generally remains a requirement. Finally, SMT solvers such as Z3 feature powerful quantifier instantiation procedures [10] (see Section 8). Considering quantifiers at the source level (as in our case) can significantly improve the performance of program verification [20].

Automated termination provers. Automated termination proving for first-order programs has long existed in industrial-grade frameworks such as ACL2 [9, 14, 21]. Recent work has introduced various procedures that tackle the higher-order case as well [11, 12, 18] through a combination of reachability analysis and upper-bounding the call-graph path lengths (size-change termination [19] in our case). Our procedure is able to automatically synthesize (and assume) inductive invariants relevant to the considered measure, thus significantly increasing its effectiveness in practice. We have found our termination prover to suffice our needs when writing interesting and applied benchmarks.

Finite model finders. All models reported by our procedure have finite nature (they have a finite range in the case of first-class functions). Techniques for finite model finding can be quite effective at counterexample reporting by explicitly considering potential models [6], uses an underlying relational constraint solver to guarantee completeness for bounded domains [22], or through incremental model search where model size is increased at each step [3, 30]. Our system uses specification execution to speed up model finding.

11. Conclusions

Our system has been implemented by gradually enriching a verifier for first-order programs with features that contribute to modular specification and verification. The resulting design aims for simplicity of the specification language: it uses the programming language expressions for specifications, and builds on well-understood concepts of contracts, invariants, and quantifiers. The system is effective when used in fully automated mode, which makes it convenient for concisely documenting and verifying algorithms and data structures.

References

- [1] L. Bachmair and H. Ganzinger. Equational reasoning in saturation-based theorem proving. In W. Bibel and P. Schmitt, editors, *Automated Deduction — A Basis for Applications*, volume I, chapter 11, pages 353–397. Kluwer, 1998.
- [2] P. Baumgartner and C. Tinelli. The Model Evolution Calculus. In F. Baader, editor, *CADE-19 – The 19th International Conference on Automated Deduction*, volume 2741 of *Lecture Notes in Artificial Intelligence*, pages 350–364. Springer, 2003.
- [3] P. Baumgartner, J. Bax, and U. Waldmann. Finite quantification in hierarchic theorem proving. In *IJCAR*, volume 8562 of *Lecture Notes in Computer Science*, pages 152–167. Springer, 2014.
- [4] C. Benzmüller and N. Sultana. LEO-II version 1.5. In *PxTP 2013*, volume 14 of *EPiC Series*, pages 2–10, 2013.
- [5] Y. Bertot and P. Castéran. *Interactive Theorem Proving and Program Development—Coq Art: The Calculus of Inductive Constructions*. Springer, 2004.
- [6] J. C. Blanchette and T. Nipkow. Nitpick: A counterexample generator for higher-order logic based on a relational model finder. In *ITP*, 2010.
- [7] S. Böhme and T. Nipkow. Sledgehammer: Judgement day. In *IJCAR*, volume 6173 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2010.
- [8] C. E. Brown. Satallax: An automatic higher-order prover. In *IJCAR*, volume 7364 of *Lecture Notes in Computer Science*, pages 111–117. Springer, 2012.
- [9] P. C. Dillinger, P. Manolios, D. Vroon, and J. S. Moore. Acl2s: "the ACL2 sedan". *Electr. Notes Theor. Comput. Sci.*, 174(2):3–18, 2007.
- [10] Y. Ge and L. M. de Moura. Complete instantiation for quantified formulas in satisfiability modulo theories. In A. Bouajjani and O. Maler, editors, *CAV*, volume 5643 of *LNCIS*, pages 306–320. Springer, 2009. doi: 10.1007/978-3-642-02658-4_25.
- [11] J. Giesl, M. Raffelsieper, P. Schneider-Kamp, S. Swiderski, and R. Thiemann. Automated termination proofs for haskell by term rewriting. *ACM Trans. Program. Lang. Syst.*, 33(2):7, 2011.
- [12] N. D. Jones and N. Bohr. Call-by-value termination in the untyped lambda-calculus. *Logical Methods in Computer Science*, 4(1), 2008.
- [13] G. Kaki and S. Jagannathan. A relational framework for higher-order shape analysis. In *ICFP*, pages 311–324. ACM, 2014.
- [14] M. Kaufmann, J. S. Moore, and P. Manolios. *Computer-aided reasoning: an approach*. Kluwer Academic Publishers, 2000.
- [15] N. Kobayashi. Types and higher-order recursion schemes for verification of higher-order programs. In Z. Shao and B. C. Pierce, editors, *POPL*, 2009.
- [16] N. Kobayashi, N. Tabuchi, and H. Unno. Higher-order multi-parameter tree transducers and recursion schemes for program verification. In *POPL*, 2010.
- [17] K. Korovin and C. Stickel. iprover-eq: An instantiation-based theorem prover with equality. In *IJCAR*, volume 6173 of *Lecture Notes in Computer Science*, pages 196–202. Springer, 2010.
- [18] T. Kuwahara, T. Terauchi, H. Unno, and N. Kobayashi. Automatic termination verification for higher-order functional programs. In *ESOP*, volume 8410 of *Lecture Notes in Computer Science*, pages 392–411. Springer, 2014.
- [19] C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The size-change principle for program termination. In *POPL*, pages 81–92. ACM, 2001.
- [20] K. R. M. Leino and C. Pit-Claudel. Trigger selection strategies to stabilize program verifiers.
- [21] P. Manolios and A. Turon. All-termination(t). In *TACAS*, volume 5505 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 2009.
- [22] A. Milicevic, J. P. Near, E. Kang, and D. Jackson. Alloy*: A general-purpose higher-order relational constraint solver. In *ICSE (1)*, pages 609–619. IEEE Computer Society, 2015.
- [23] P. C. Nguyen and D. V. Horn. Relatively complete counterexamples for higher-order programs. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 446–456, 2015.
- [24] P. C. Nguyen, S. Tobin-Hochstadt, and D. V. Horn. Higher-order symbolic execution for contract verification and refutation. *CoRR*, abs/1507.04817, 2015.
- [25] R. Nieuwenhuis and A. Rubio. Paramodulation-based theorem proving. In *Handbook of Automated Reasoning (Volume 1)*, chapter 7. Elsevier and The MIT Press, 2001.
- [26] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*. Springer, 2002.
- [27] C. L. Ong and S. J. Ramsay. Verifying higher-order functional programs with pattern-matching algebraic data types. In *POPL*, 2011.
- [28] B. Pierce. *Types and Programming Languages*. The MIT Press, Cambridge, Mass., 2001.
- [29] V. Prevosto and U. Waldmann. SPASS+T. In *ESCoR: Empirically Successful Computerized Reasoning*, volume 192, 2006.
- [30] A. Reynolds, C. Tinelli, A. Goel, and S. Krstic. Finite model finding in SMT. In *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 640–655. Springer, 2013.
- [31] J. C. Reynolds. *Logical foundations of functional programming*, chapter Introduction. Addison-Wesley, 1990.
- [32] P. M. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In *PLDI*, 2008.
- [33] N. Sultana, J. C. Blanchette, and L. C. Paulson. LEO-II and satallax on the sledgehammer test bench. *J. Applied Logic*, 11(1):91–102, 2013.
- [34] P. Suter, M. Dotta, and V. Kuncak. Decision procedures for algebraic data types with abstractions. In *POPL*, 2010.
- [35] P. Suter, A. S. Köksal, and V. Kuncak. Satisfiability modulo recursive programs. In *SAS*, 2011.
- [36] H. Unno, T. Terauchi, and N. Kobayashi. Automating relatively complete verification of higher-order functional programs. In *POPL*, pages 75–86. ACM, 2013.
- [37] N. Vazou, P. M. Rondon, and R. Jhala. Abstract refinement types. In *ESOP*, volume 7792 of *Lecture Notes in Computer Science*, pages 209–228. Springer, 2013.
- [38] N. Vazou, E. L. Seidel, and R. Jhala. Liquidhaskell: experience with refinement types in the real world. In *Haskell*, pages 39–51. ACM, 2014.
- [39] N. Vazou, E. L. Seidel, R. Jhala, D. Vytiniotis, and S. L. P. Jones. Refinement types for haskell. In *ICFP*, 2014.
- [40] N. Voirol, E. Kneuss, and V. Kuncak. Counter-example complete verification for higher-order functions. In P. Haller and H. Miller, editors, *Scala@PLDI*, pages 18–29. ACM, 2015. doi: 10.1145/2774975.2774978.

APPENDIX

A. Syntax of PureScala

B. Encoding Pattern Match Expressions

Consider pattern matching expressions that deconstruct algebraic data type instances into their component fields. We require pattern matches to cover all ADT constructors. Given a set of n cases where pattern match case $i \in \{1, \dots, n\}$ has shape

$$\text{case}_i ::= \text{case cons}_i(v_{i,1}, \dots, v_{i,m}) \Rightarrow E_i$$

and given pattern match expression

$$E ::= \text{SCRUT match } \{ \text{case}_1 \dots \text{case}_n \},$$

the transformation of (b, E) closely resembles that of **if**-expressions. For each case_i , we define b_i a fresh boolean-sorted constant. Given $E : T$, we also define the fresh constant r with sort $\mathcal{S}(T)$. Given $(b, \text{SCRUT}) \rightsquigarrow^t s$, we consider for each case_i the substitution $\theta_i = \{ \mathcal{V}(v_{i,j}) \rightarrow \text{field}_{i,j}(s) \mid 1 \leq j \leq m \}$ and compute

$(b_i, E_i) \rightsquigarrow^t e_i$ for $1 \leq i \leq n$. This lets us define the clausal encoding of pattern match expressions as

$$\phi_{split} = \{ (b \wedge isCons_i(s)) \iff b_i \mid 1 \leq i \leq n \} \\ \cup \{ b_i \implies (r \simeq \theta_i[e_i]) \mid 1 \leq i \leq n \}.$$

Note that unlike in the **if**-expression case for b_t and b_e , we don't need to explicitly ensure $b_i \implies \neg b_j$ when $i \neq j$ as this is ensured by the ADT theory ($isCons_i(s)$ is mutually exclusive with $isCons_j(s)$). This leads to our final inference rule

$$\text{MATCH} \frac{r, b_1, \dots, b_n \text{ fresh} \quad (b, \text{SCRUT}) \rightsquigarrow (s, \phi_s, A_s, F_s) \\ (b_i, E_i) \rightsquigarrow (e_i, \phi_i, A_i, F_i), 1 \leq i \leq n}{(b, E) \rightsquigarrow (r, \\ \phi_s \cup \theta_1[\phi_1] \cup \dots \cup \theta_n[\phi_n] \cup \phi_{split}, \\ A_s \cup \theta_1[A_1] \cup \dots \cup \theta_n[A_n], \\ F_s \cup \theta_1[F_1] \cup \dots \cup \theta_n[F_n])}$$

We use $\theta_i[\cdot]$ on all terms stemming from the transformation of E_i to ensure correspondence between the match deconstruction semantics and the ADT theory axiomatization. Our transformation therefore guarantees that $field_{i,j}(s)$ terms can only occur in positions where $isCons_i(s)$ must hold.

C. Proofs

Lemma 1. *There exists model $M \models b \wedge \phi_e \wedge \phi_v \wedge e \simeq v$.*

Proof. First note that v is not necessarily ground as it may contain constants introduced by the transformation of first-class functions in v . However, ϕ_v must be empty given the definition of value expression v . Furthermore, given two expressions E_1, E_2 and constant b with $(b, E_i) \rightsquigarrow (e_i, \phi_{E_i}, -, -)$, $m[E_i] \xrightarrow{*} v_i$, and $(b, v_i) \rightsquigarrow^t v_i$ for $i \in \{1, 2\}$, we have $vars(v_1) \cap vars(v_2) = \emptyset$ as all introduced constants are fresh. Therefore, given models $M_i \models b \wedge \phi_{E_i} \wedge e_i \simeq v_i, i \in \{1, 2\}$, there must exist a model $M \models b \wedge \phi_{E_1} \wedge \phi_{E_2} \wedge e_1 \simeq v_1 \wedge e_2 \simeq v_2$. It is then trivial to conclude the proof by induction over the structure of E . \square

Lemma 3. *There exists an M satisfying Lemma 1 such that for $E_C \subseteq E$ where $E_C ::= C(E_1, \dots, E_n)$ and for $fun \in F_E$, if $M \models C_b(E_C) \wedge (fun)_b$, then $M \models C_i(C) \simeq (fun)_f$ iff $C_v(C) = C_v(C_E((fun)_f))$.*

Proof. Follows directly from Lemma 1 and Proposition 2. Indeed, if $C_t(C) \simeq (fun)_f$ must hold, then C and $C_E((fun)_f)$ are one and

```

definition ::= abstract class id tdecls
              | case class id tdecls (decls) extends id tdecls
              | def id tdecls (decls) : type = expr
tdecls ::=  $\epsilon \mid [id \langle, id \rangle^*]$ 
decls ::=  $\epsilon \mid id : type \langle, id : type \rangle^*$ 
expr ::= true | false | id
        | if (expr) expr else expr
        | id tparams (  $\langle expr \langle, expr \rangle^* \rangle?$  )
        | expr (  $\langle expr \langle, expr \rangle^* \rangle?$  )
        | (decls)  $\Rightarrow$  expr
        | expr match {
              < case id (  $\langle id \langle, id \rangle^* \rangle?$  )  $\Rightarrow$  expr  $\rangle^+$ 
            }
tparams ::=  $\epsilon \mid [type \langle, type \rangle^*]$ 
type ::= id tparams | Boolean
id ::= IDENT

```

Figure 6. Syntax of PureScala

the same in $m[E]$. Otherwise, an arbitrary (different) value can be chosen for the constant $(fun)_f$. \square

Lemma 4. *Given inputs m with $m[P] \xrightarrow{*} v$ and $(b_p, v) \rightsquigarrow^t v_p$, given $E_C \subseteq P$ where $C_t(E_C) = \text{dispatch}_{type}(c, a_1, \dots, a_n)$ and $(C_b(E_C), C_v(E_C)) \rightsquigarrow^t v_C$, there exists model M such that either*

1. $M \models b_p \wedge \phi_p \wedge e_p \simeq v_p \wedge \theta_f[\phi_f] \wedge \theta_f[e_f] \simeq v_C$, or
2. $M \models b_p \wedge \phi_p \wedge e_p \simeq v_p \wedge \neg b_f$.

Proof. Follows from Lemma 3 and the correspondence between the definition of the θ_f substitution and the operational semantics for both functions (closure substitution) and function applications (argument substitution). \square

Theorem 5. *Given $P : \text{Boolean}$ with ϕ_i from $U(P)$, if $\phi_i \in \text{Unsat}$, then no inputs m exist such that $m[P] \xrightarrow{*} \text{false}$.*

Proof. Let us assume inputs m exist such that $m[P] \xrightarrow{*} \text{false}$. Lemma 1 tells us that there exists $M_0 \models b_p \wedge \neg e_p \wedge \phi_p$. Lemma 4 then gives us a model $M_1 \models \phi_1$ and the proof follows through by induction on M_i . \square

Lemma 6. *Given ϕ_i from $U_{\text{POST}}(E_{\text{POST}})$, if $\phi_i \in \text{Unsat}$ and POST terminates on all inputs, then there exist no inputs m such that $m[\text{POST}] \xrightarrow{*} \text{false}$.*

Proof. First off, evaluation semantics of function application clearly ensure a correspondance between inputs to POST and E_{POST} . Let us assume there exist inputs m such that $m[\text{POST}] \xrightarrow{*} \text{false}$. Consider the set S defined as the smallest fixpoint where

1. $\mathbf{f}(m[v_1], \dots, m[v_n]) \in S$,
2. if $\mathbf{f}(E_1, \dots, E_n) \in S$, then for each $\mathbf{f}(E'_1, \dots, E'_n)$ encountered during evaluation of $\{v^n \mapsto E^n\}[\text{POST}]$, $\mathbf{f}(E'_1, \dots, E'_n) \in S$.

We assume m is selected such that for each $\mathbf{f}(E_1, \dots, E_n) \in S$, either $E_i = m[v_i], 1 \leq i \leq n$, or $\{v^n \mapsto E^n\}[\text{POST}] \xrightarrow{*} \text{true}$. Note that such inputs can be obtained by recursively selecting the arguments of the offending application in S as inputs m . Further note that termination of each $\{v^n \mapsto E^n\}[\text{POST}]$ is ensured by the condition on POST .

Lemma 1 ensures that for each $\mathbf{f}(E_1, \dots, E_n) \in S$ with associated $b_f, \text{app}_j, \text{fun}_j$ where $(\text{fun}_j)_f = \mathbf{f}$, given the clause set obtained by $(\phi_{\text{POST},j}, -, -) = \mathcal{P}_{\text{POST}}(\text{app}_j, \text{fun}_j)$, there exists a model $M_j \models b_f \wedge \phi_{\text{POST},j}$. Finally, we can adapt option No. 1 of Lemma 4 to $M \models b_p \wedge \phi_p \wedge \neg e_p \wedge \theta_f[\phi_p] \wedge \theta_f[p]$, thus concluding our proof. \square

Theorem 7. *Given ϕ_i from $U_{\text{POST}}(P)$, if $\phi_i \in \text{Unsat}$, POST is valid, and POST terminate on all inputs, then there exist no inputs m such that $m[P] \xrightarrow{*} \text{false}$.*

Proof. We assume inputs m exist such that $m[P] \xrightarrow{*} \text{false}$. Validity and termination of POST , along with Lemma 6, ensure that there exists $M_j \models \phi_{\text{POST},j}$ for each $(\phi_{\text{POST},j}, -, -) = \mathcal{P}_{\text{POST}}(\text{app}_j, \text{fun}_j)$. Theorem 5 then ensures there exists $M_p \models \phi_{p,k}$ from $U(P)$ and Lemma 4 finally guarantees that M_p and all M_j can be unified into a single model $M \models \phi_i$ from $U_{\text{POST}}(P)$. \square

Theorem 11. *If for each $E \in \bigcup_{\mathbf{f} \in P} C_f$ there exists some ϕ_i from $U(\text{DEC}(E))$ such that $\phi_i \in \text{Unsat}$, then P terminates for all inputs.*

Proof. Given $E \in C_f$, Theorem 5 ensures that no inputs m exist such that $m[\text{DEC}(E)] \xrightarrow{*} \text{false}$, namely, either $m[\text{DEC}(E)] \xrightarrow{*} \text{true}$ or $m[\text{DEC}(E)]$ diverges. However, if $m[\text{DEC}(E)]$ does not terminate, then there must exist an infinite sequence of function application evaluations $\text{APP}_0, \text{APP}_1, \dots$ such that $\text{APP}_0 ::= \mathbf{f}(m[\mathbf{f}.args])$

and $(APP_i, APP_{i+1}) \in R$. For each $APP_i ::= \mathbf{f}_{i+1}(A_{i,1}, \dots, A_{i,n_i})$, there exists by definition some $\mathbf{f}_{i+1}(E_{i,1}, \dots, E_{i,n_i}) \in \mathcal{C}_{\mathbf{f}_i}$ and inputs m_i such that $m_i \llbracket \text{path}(\mathbf{f}_{i+1}(E_{i,1}, \dots, E_{i,n_i})) \rrbracket \xrightarrow{*} \mathbf{true}$ and $m_i \llbracket E_{i,j} \rrbracket \xrightarrow{*} A_{i,j}$ for $1 \leq j \leq n_i$. Given $\sum_j \text{size}(A_{i,j}) \xrightarrow{*} v_i$, we therefore have an infinite sequence $v_1 > v_2 > \dots \geq 0$ and obtain a contradiction. Applying Proposition 10 then concludes our proof. \square

Theorem 13. For $\mathbf{f} \in \mathcal{P}$, if $\text{Posts}(\mathbf{f})$ is post-terminating for all inputs, $\text{Posts}(\mathbf{f})$ are valid, and for each $E \in \bigcup_{g \in \mathcal{P}} C_g$ there exists some ϕ_i from $U_{\text{POST}}(\text{Posts}(\mathbf{f}), \text{DEC}(E))$ such that $\phi_i \in \text{Unsat}$, then \mathcal{P} terminates for all inputs.

Proof. Follows from conditions of U_{POST} soundness begin satisfied and similar arguments to Theorem 11. \square

Lemma 16. For $fun_1, fun_2 \in F$ where $fun_j = (_, f_j, \mathbf{f}_j, _)$ and \mathbf{f}_j is ground for $j \in \{1, 2\}$, we have $\mathbf{f}_1 = \mathbf{f}_2$ iff some model $M \models f_1 \simeq f_2 \wedge EQ_{\mathbf{f}}(F)$ exists.

Proof. Proposition 15 ensures that there is an $M \models EQ_{\mathbf{f}}(F)$. By induction on the definition of eq , one sees that the only non-trivial instance is case No. 1 when $C_t(E_1)$ or $C_t(E_2)$ contains some term with sort $\sigma_{\mathbf{f}}$. However, as the \mathbf{f}_j are ground, the definition of $EQ_{\mathbf{f}}$ ensures that there is a $x \simeq y$ clause in $EQ_{\mathbf{f}}(F)$ for such terms. One then concludes the proof by induction on the set of first-class functions in \mathbf{f}_j that satisfy the conditions of case No. 1. \square

Lemma 18. Given a model $M \models \eta_i(v)$, there is no term $s \subset M \llbracket v \rrbracket$ with sort $\sigma_{\mathbf{f}}$ such that $M \models s \notin E_i(v)$.

Proof. Trivial if $\sigma = \sigma_{\mathbf{f}}$. The clause $\eta_i(v)$ ensures there exists $t_j, j \leq i$ such that $M \models v \simeq t_j$. The $\sigma_{\mathbf{f}}$ -exhaustiveness constraint on the enumeration of σ ensures that there exists a model M_2 and term t_k such that $M_2 \models M \llbracket v \rrbracket \simeq t_k$ and for any $s \subset M \llbracket v \rrbracket$ with sort $\sigma_{\mathbf{f}}$, there exists $s_k \subset t_k$ such that $M_2 \models s \simeq s_k$. Unicity further guarantees that $j = k$ and therefore $M \llbracket v \rrbracket = M_2 \llbracket t_k \rrbracket$. \square

Theorem 20. Given a model $M \models \phi_i \cup \psi_i \cup \rho_i$, we have $\mathcal{C}_M(\text{vars}(\mathcal{P})) \llbracket \mathcal{P} \rrbracket \xrightarrow{*} \mathbf{false}$.

Proof. We start by defining \mathcal{C}'_M , an extension to \mathcal{C}_M with the following extra rule

$$\frac{fun \in F_i, fun = (_, f, \mathbf{f}, c^n) \quad E_i = \mathcal{C}'_M(c_i, \mathbf{f}.closures_i.tpe), 1 \leq i \leq n}{\mathcal{C}'_M(t, \mathbf{f}.tpe) = \{ \mathbf{f}.closures \mapsto E^n \} \llbracket \mathbf{f} \rrbracket}$$

It is clear given the definition of \mathcal{C}'_M that there exists a left inverse to \mathcal{C}'_M , which we will call \mathcal{C}''_M , such that $\mathcal{C}''_M(\mathcal{C}'_M(t, \mathcal{T})) = t$.

Consider the constant-expression pair b, E that was transformed during $U(\mathcal{P})$ and consider the corresponding $(b, E) \rightsquigarrow (e, \phi, A, F)$. Note that $\text{vars}(E) - \text{vars}(\mathcal{P})$ is not necessarily empty due to term substitutions in pattern-matching expression transformation and $\mathcal{I}(app_i, fun_i)$ dispatch. Although (e, ϕ, A, F) may appear substituted in $U(\mathcal{P})$, it is clear that one can determine the subset of app_j, fun_j pairs in D_i that are related to E . We therefore let Φ_E consist of the unsubstituted portion of $\phi_i \cup \psi_i \cup \rho_i$ that stems from E (note that $\phi \subseteq \Phi_E$). We show by induction that if there exists $M_E \models \{b\} \cup \Phi_E$, then $\mathcal{C}_{M_E}(\text{vars}(E)) \llbracket E \rrbracket \xrightarrow{*} \mathcal{C}'_{M_E}(e, E.tpe)$.

Let us consider the case of pattern-matching expressions where

$$E ::= \mathbf{S} \text{ match } \{ \dots \text{ case } \text{cons}_j(v_{j,1}, \dots, v_{j,m_j}) \Rightarrow R_j \dots \}.$$

By induction, we have $\mathcal{C}_{M_E}(\text{vars}(S)) \llbracket S \rrbracket \xrightarrow{*} v$ and can assume without loss of generality that $v ::= \text{cons}_j(E_1, \dots, E_{m_j})$. Consider the model $M_{R_j} = M_E \{ \mathcal{V}(v_{j,k}) \mapsto \mathcal{C}''_{M_E}(E_k) \mid 1 \leq k \leq m_j \}$ and

note that $M_E \models \{b\} \cup \Phi_E$ implies that $M_{R_j} \models \{C_b(R_j)\} \cup \Phi_{R_j}$ and $\mathcal{C}'_{M_E}(e, E.tpe) = \mathcal{C}'_{M_{R_j}}(C_t(R_j), R_j.tpe)$.

Now consider the case of function application expressions where $E ::= C(E_1, \dots, E_n)$, and consider $(C_b(E), E) \rightsquigarrow \text{disp}(app)$ where $app = (_, c, _, a^n)$. We further consider the unsubstituted sets A_E, F_E and $T_E(app)$ that stem from E . We examine two mutually exclusive cases:

$M_E \models c \in T_E(app)$: consider $fun \in F_E$ corresponding to c where $fun = (_, _, \mathbf{f}, c^m)$, as well as $C_j = \mathcal{C}'_{M_E}(c_j, v_j.tpe)$ for $v_j \in \mathbf{f}.closures$ and $A_k = \mathcal{C}'_{M_E}(a_k, v_k.tpe)$ for $v_k \in \mathbf{f}.args$. Let us now define the following models

$$M_{\mathbf{f}} = M_E \{ \mathcal{V}(v_j) \mapsto \mathcal{C}''_{M_E}(C_j) \mid v_j \in \mathbf{f}.closures \},$$

$$M_{\mathbf{f}.body} = M_{\mathbf{f}} \{ \mathcal{V}(v_k) \mapsto \mathcal{C}''_{M_E}(A_k) \mid v_k \in \mathbf{f}.args \}.$$

Since $M_E \models \{c \in T_E(app), C_b(E)\} \cup \Phi_E$, we have that $M_{\mathbf{f}.body} \models \{b_{\mathbf{f}}\} \cup \Phi_{\mathbf{f}.body}$. Given $(b_{\mathbf{f}}, \mathbf{f}.body) \rightsquigarrow \text{body}$, we therefore have $\mathcal{C}'_{M_E}(e, E.tpe) = \mathcal{C}'_{M_{\mathbf{f}.body}}(body, \mathbf{f}.body.tpe)$. Finally, Lemma 16 ensures consistency of the uninterpreted function symbol applications $\text{disp}(app)$ given the ground first-class function $\{v_j \mapsto C_j \mid v_j \in \mathbf{f}.closures\} \llbracket \mathbf{f} \rrbracket$.

$M_E \models c \in E_{all,i}(\text{vars}(\mathcal{P}), A_E)$: the ψ_i clause set ensures that $M \llbracket E_{all,i}(\text{vars}(\mathcal{P}), A_E) \rrbracket \cap M \llbracket T_E(app) \rrbracket = \emptyset$, and therefore $\mathcal{C}'_{M_E}(c, C.tpe) = \mathcal{C}_{M_E}(c, C.tpe)$. To conclude this case, it suffices to note that $\mathcal{C}_{M_E}(\text{vars}(E)) \llbracket E \rrbracket \xrightarrow{*} \mathcal{C}'_{M_E}(e, E.tpe)$ by definition of \mathcal{C}_{M_E} .

The remaining cases for E follow trivially from the inductive hypothesis. To conclude our proof, one notes that $M \models \{b_{\mathcal{P}}\} \cup \Phi_{\mathcal{P}}$ where $\Phi_{\mathcal{P}} = \phi_i \cup \psi_i \cup \rho_i$ and $M \models \neg C_t(\mathcal{P})$, therefore we have $\mathcal{C}'_M(C_t(\mathcal{P})) = \mathbf{false}$ and $\mathcal{C}_M(\text{vars}(\mathcal{P})) \llbracket \mathcal{P} \rrbracket \xrightarrow{*} \mathbf{false}$. \square

Theorem 21. If inputs m exist such that $m \llbracket \mathcal{P} \rrbracket \xrightarrow{*} \mathbf{false}$, then there exists $i \in \mathbb{N}$ and model M such that $M \models \phi_i \cup \psi_i \cup \rho_i$.

Proof. By Proposition 15, Proposition 17 and Proposition 19 along with freshness of all relevant constants, we have $M_{\psi} \models \psi_i$ existence. Theorem 5 further ensures that $M_{\phi} \models \phi_i$ exists and the proof of Theorem 5 in combination with Lemma 16 ensures that $M_{\phi, \psi} \models \phi_i \cup \psi_i$ exists as well.

Let us now consider ρ_i . We define the set A_m of all app tuples corresponding to function applications encountered during evaluation of $m \llbracket \mathcal{P} \rrbracket$. Similarly, let F_m the set of fun tuples corresponding to encountered first-class functions that are defined within \mathcal{P} . Finally, let F_m the set of first-class functions defined in m that can appear in caller position of $app \in A_m$. Note that this restriction filters out the first-class functions in equality positions of \mathcal{C}_M outputs.

Our fairness requirement ensures that $(app)_c \in T_i(app)$ will eventually be satisfied for $app \in A_m$ when $(app)_c$ corresponds to some $fun \in F_m$. For $\mathbf{f} \in F_m$, consider \mathbf{f}_2 of the shape returned by \mathcal{C}_M such that \mathbf{f} and \mathbf{f}_2 agree on all relevant inputs given by A_m and the shape of \mathbf{f}_2 ensures distinctness from functions in F_m (ensured by definition of \mathcal{C}_M). As no equality predicate on function types exists in \mathcal{P} , given $m_2 = \{ \mathbf{f} \mapsto \mathbf{f}_2 \} \llbracket m \rrbracket$, it is clear that $m_2 \llbracket \mathcal{P} \rrbracket \xrightarrow{*} \mathbf{false}$. There must therefore exist some inputs m_k such that $m_k \llbracket \mathcal{P} \rrbracket \xrightarrow{*} \mathbf{false}$ and the associated F_{m_k} contains only functions of the shape returned by \mathcal{C}_M , thus ensuring satisfiability of $\phi_i \cup \psi_i \cup \rho_i$ for some $i \in \mathbb{N}$. \square

Lemma 23. Given a model $M \models \eta_i(v)$, there is no term $s \subset M \llbracket v \rrbracket$ with sort $\sigma_{\mathcal{T}}$ such that $M \models s \notin S_i(v)$.

Proof. Consider the case where $\sigma = \sigma_{\mathcal{T}}$. Unlike for Lemma 18, this case is non-trivial as there may exist $s \subset M \llbracket v \rrbracket$ with sort $\sigma_{\mathcal{T}}$. Furthermore, given t_j such that $M \models v \simeq t_j$, theory separation doesn't ensure that there is a corresponding $s_j \subset t_j$ such that

$s = M \llbracket s_j \rrbracket$ as t_j and s have the same sort σ_{\top} . However, ADT theory axioms do ensure that such a correspondance exists, an observation that is not clear for any arbitrary theory. The rest of the proof closely follows the general case from Lemma 18. \square

Theorem 25. For $(\phi_i, A_i, F_i) \in U_{\text{INV}}(\mathbb{P})$ with associated ψ_i, ρ_i , if $M \models \phi_i \cup \psi_i \cup \rho_i$, then INV holds on $\mathcal{C}_M(\text{vars}(\mathbb{P}))$.

Proof. Given $v \in \text{vars}(\mathbb{P})$, it follows from Lemma 23 and induction on \triangleright that for each $E \subseteq \mathcal{C}_M(v)$ such that $E : \top$, there exists $s \in S_{\text{all},i}(\text{vars}(\mathbb{P}), A_i)$ such that $\mathcal{C}_M(s, \top) = E$. Furthermore, the b_{\top} constant corresponding to s must hold in M .

Consider the subsets of $\phi_i \cup \psi_i \cup \rho_i$, A_i and F_i that derive from $\mathcal{P}_{\text{INV}}(s)$ and note that they correspond to the sequence $U(\neg \text{INV})$ where $\mathcal{V}(x)$ has been substituted by s . We can therefore construct a model $M_s = M \{ \mathcal{V}(x) \mapsto M \llbracket s \rrbracket \}$ such that for some $(\phi_j, A_j, F_j) \in U(\neg \text{INV})$ with associated ψ_j, ρ_j , we have $M_s \models \phi_j \cup \psi_j \cup \rho_j$. The proof of Theorem 20 then ensures that $\mathcal{C}_{M_s}(\text{vars}(\text{INV})) \llbracket \text{INV} \rrbracket \xrightarrow{*} \text{true}$.

Finally, as $x \in \text{vars}(\mathbb{P})$ when an instance v_{\top} is required for $\mathcal{C}_M(\text{vars}(\mathbb{P}))$, $\mathcal{C}_M(x)$ will produce a value on which INV holds. Note that when no value v_{\top} exists, either M doesn't exist, or $\mathcal{C}_M(x)$ won't terminate, thus ensuring soundness of produced inputs. \square

Theorem 26. If inputs m exist such that INV holds on m and $m \llbracket \mathbb{P} \rrbracket \xrightarrow{*} \text{false}$, then there exists $(\phi_i, A_i, F_i) \in U_{\text{INV}}(\mathbb{P})$ with associated ψ_i, ρ_i such that $M \models \phi_i \cup \psi_i \cup \rho_i$ exists.

Proof. Closely follows the proof of Theorem 21 where m is converted into m_k such that all first-class functions in m_k correspond to results of \mathcal{C}_M . Note that as INV holds on m , any v_{res} that depends on \top can be obtained by evaluating the corresponding first-class function in m at an arbitrary point. \square

Lemma 27. For $Q(A_p) = (\mathcal{X}_0, _), (\mathcal{X}_1, _), \dots, \bigcup_i \mathcal{X}_i = \mathcal{X}(A_p)$.

Proof. Equivalence for INIT-G and INIT-X rules is given by definition. For EXPAND rule to apply, there must exist $app \in A_0$ such that app is not quantifier-free. The EXPAND rule application then corresponds exactly to introducing $(b \wedge b_s, \theta_s \llbracket c \rrbracket, tpe, \theta_s \llbracket a^n \rrbracket)$ into A_1 for each $app \in A_0$ that is not quantifier-free and applying INIT-G. The proof then follows by induction on i . \square

Theorem 28. Given $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$, if $\phi_i \in \text{Unsat}$, then there exist no inputs m on which $\forall x_1, \dots, x_m. \text{PROP}$ holds.

Proof. Follows from the Compactness Theorem as each $\mathcal{P}_{\text{QUANT}}(x_i)$ corresponds to a quantifier-free instantiation of PROP. \square

Lemma 29. For $(\mathcal{X}_i, A_i) \in Q(A_p)$ and model M , if M finitely interprets \mathcal{X}_i , then M finitely interprets \mathcal{X}_{i+1} .

Proof. Condition No. 1 clearly carries over to all $x \in \mathcal{X}_{i+1} - \mathcal{X}_i$. If condition No. 2 holds on $x_{i+1} \in \mathcal{X}_i$ selected for A_{i+1} , then given $(b_{i+1}, s_{i+1}) \in X_k(A_{i+1})$, there exists $(b_i, s_i) \in X_k(A_i)$ such that $M \models b_{i+1} \iff b_i$ and $M \models s_{i+1} \simeq s_i$. Therefore, all $x \in \mathcal{X}_{i+1} - \mathcal{X}_i$ will satisfy condition No. 2. \square

Theorem 30. Given $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$ with associated ψ_i, ρ_i, γ_i , and $M \models \phi_i \cup \psi_i \cup \rho_i \cup \gamma_i$, if $\forall x_1, \dots, x_m. \text{PROP}$ is in a semi-decidable fragment, then $\forall x_1, \dots, x_m. \text{PROP}$ holds on $\mathcal{C}_M(\text{vars}(\forall x_1, \dots, x_m. \text{PROP}))$.

Proof. Let us assume there exist no inputs m such that $\forall x_1, \dots, x_m. \text{PROP}$ holds on m . Semi-decidability then tells us there exists $j > i$ such that $\phi_j \cup \psi_j \cup \rho_j \cup \gamma_j \in \text{Unsat}$.

Let us consider u_{i+1} and $\Phi_{i+1} = \phi_{i+1} \cup \psi_{i+1} \cup \rho_{i+1} \cup \gamma_{i+1}$. We start by assuming u_{i+1} was obtained through \mathcal{I} . Note that if

$$M \not\models (app_{i+1})_b \wedge (fun_{i+1})_b \wedge (app_{i+1})_f \simeq (fun_{i+1})_f,$$

then all clauses introduced by \mathcal{I} are irrelevant and the tuples in $A_{i+1} - A_i$ satisfy condition No. 1 of *finite interpretation*. Otherwise, $M \models \rho_i$ ensures that a corresponding inlining has already taken place and there must therefore exist a mapping θ from fresh variables introduced by \mathcal{I} to variables in M such that $\theta \llbracket u_{i+1} \rrbracket \subseteq u_i$ and $M \{ x \mapsto M \llbracket y \rrbracket \mid x \mapsto y \in \theta \} \models \Phi_{i+1}$.

If u_{i+1} was obtained through $\mathcal{P}_{\text{QUANT}}$, then $\phi_{i+1}, \psi_{i+1}, \rho_{i+1}$ are unchanged and Lemma 29 ensures that there exists model $M_{i+1} \models \Phi_{i+1}$. Finally, by induction on i we have existence of $M_k \models \Phi_k$ for all $k > i$ and thus obtain a contradiction. Validity of $\mathcal{C}_M(\text{vars}(\forall x_1, \dots, x_m. \text{PROP}))$ is then ensured by applying the proof of Theorem 20. \square

Theorem 31. If there exist inputs m such that for each first-class function $f \in m$, the set $\{ v \mid f(E_1, \dots, E_n) \xrightarrow{*} v \}$ is finite and $\forall x_1, \dots, x_m. \text{PROP}$ is in a semi-decidable fragment, then there exists $(\phi_i, A_i, F_i) \in U_{\text{QUANT}}(\text{PROP})$ with associated ψ_i, ρ_i, γ_i , and model M such that $M \models \phi_i \cup \psi_i \cup \rho_i \cup \gamma_i$.

Proof. Follows from semi-decidability and fair exploration of relevant arguments to each f . \square

D. Examples

D.1 Lambda Calculus Evaluator

```

object LambdaCalculus {
  abstract class Term
  case class Var(x: BigInt) extends Term
  case class Abs(x: BigInt, body: Term) extends Term
  case class App(func: Term, arg: Term) extends Term

  def fv(t: Term): Set[BIGInt] = t match {
    case Var(x) => Set(x)
    case Abs(x, body) => fv(body) ++ Set(x)
    case App(func, arg) => fv(func) ++ fv(arg)
  }

  // [x->u]t
  def subst(x: BigInt, u: Term, t: Term): Term = t match {
    case Var(y) => if (x == y) u else t
    case Abs(y, body) => if (x == y) t else Abs(y, subst(x, u, body))
    case App(f, a) => App(subst(x, u, f), subst(x, u, a))
  }

  // big step call-by-value evaluation
  def eval(t: Term): Option[Term] = t match {
    case App(t1, t2) => eval(t1) match {
      case Some(Abs(x, body)) => eval(t2) match {
        case Some(v2) => eval(subst(x, v2, body))
        case None() => None[Term]()
      }
      case _ => None[Term]() // stuck
    }
    case _ => Some(t) // Abs or Var, already a value
  }
  ensuring { res => res match {
    case Some(t) => isValue(t)
    case None() => true
  } }

  def isValue(t: Term): Boolean = t match {
    case Var(x) => true
    case Abs(x, body) => true
  }
}

```



```
    case App(f, a) => false
  }
}
/* Termination checker (LoopProcessor) says:
   Non-terminating for call:
   eval(App(Abs(0, App(Var(0), Var(0))), Abs(0, App(Var(0), Var(0))))))
*/
```